MARKED-UP VERSION OF SUBSTITUTE SPECIFICATION UNDER 37 C.F.R. 1.125

Per Zone Variable BPI for Improving Storage Device Capacity and Yield

By

Saeed Asgari

Mathew Vea

And

George Iszlai

Field of the Invention

The present invention relates ~~generally~~ to ~~the storage of~~ information storage on a ~~fixed~~ storage media~~, and more particularly to improving storage of information on rotating magnetic media~~ such as a disk~~s~~ in a disk drive.

Background of the Invention

Data storage devices such as disk drives are used in many data processing systems~~ for data storage~~. Typically a disk drive includes a magnetic data disk having disk~~recording~~ surfaces with concentric data tracks, and a transducer head paired with each disk~~recording~~ surface~~,~~ for reading data from and writing data to~~, and reading data from,~~ the data tracks.~~ Each paired magnetic head and media surface couples to provide a unique data recording capability which depends on the fly height of the head from the recording surface, the quality/distribution of magnetic media on the recording surface, and the magnetic properties of the magnetic head.~~

Disk drive storage capacity increases by increasing the data density (or areal density) of the data stored on the disk surfaces. Data density is the linear bit density on the tracks multiplied by the track density across the disk surface. Data density is measured in bits per square inch (BPSI), linear bit density is measured in bits per inch (BPI) and track density is measured in tracks per inch

(TPI). As data density increases, the head performance distribution also increases which diminishes disk drive storage capacity and yield.

C~~C~~onventional disk drives fail to account for ~~methods of recording data using the paired head and recording surface are inefficient because they do not take into consideration~~ the different~~ces in~~ ~~data recording~~ capabilities of~~between~~ the ~~one pair of~~ head and disk~~recording~~ surface pairs~~, and another head and surface pair~~. Conventionally, each disk surface is formatted to store the same amount of data as every other disk surface. However, each head and disk surface pair has unique data recording capability, such as sensitivity and accuracy, which depends on the fly height of the head over the disk surface, the magnetic properties of the head and the quality/distribution of the magnetic media for the disk surface. Thus, in conventional disk drives a head and disk surface pair that has a low error rate is formatted to the same BPI and TPI as a head and disk surface pair that has a high error rate.

~~Though the heads are designed to perform identically in read/write operations, in practice different heads in a disk drive can have different read/write performance capabilities. Lower performing heads cannot read/write data as that of other heads in the disk drive. Typically, a single error rate level and a single storage capacity level are used to record data for all the pair heads and surfaces. This results in inefficient data storage for those pairs of heads and surfaces that can store more data. It also lowers the qualification yields of the disk drives because one or more pairs of heads and surfaces do not record data at the qualifying error rate and capacity levels.~~

~~Further, in high data rate design of disk drives, as the recording density (i.e. bits-per-inch and/or tracks-per-inch) is increased, maintaining transducer head tolerances has become a challenge. Variance in the relative head performance distribution increases with increasing data density. In conventional~~

~~disk drives, the drive yield and capacity suffers as a result of head performance variations in disk drives.~~

~~One method of increasing the data storage capacity of a disk drive includes increasing the areal density of the data stored on the media surfaces (bits/sq. in. -- BPSI). Areal density is the track density which is the number of tracks per radial inch (TPI) that can be packed onto the media/recording surface, multiplied by the linear density (BPI) which is the number of bits of data that can be stored per linear inch.~~

Conventional disk drive manufacturing~~processes~~ applies a single error rate and a single data storage level for the head and disk surface pairs, and ~~for qualifying disk drives~~ scraps~~ a~~ disk drives that include a low performing head and disk surface pair that fails to meet ~~when~~ the qualifying requirements~~ measured disk capacity of the disk drive is less than a target disk capacity~~. This lowers storage capacity due to inefficient use of high performing head and disk surface pairs that can store more data, and lowers yield due to disk drives being scrapped if they include a low performing head and disk surface pair even if they also include a high performing head and disk surface pair~~Conventionally, each recording surface is formatted to store the same amount of data as every other recording surface. Thus, a recording surface that has a low error rate is formatted to the same TPI and BPI levels, as a recording surface having a high error rate, even though it can store more data. However, by adopting a single TPI and BPI level for every recording surface, conventional processed fail to account for the differences in sensitivity and accuracy of the paired head and recording surface, which results in less data storage and more waste of space on each recording surface. This also results in lower overall yields of disk drives because if even a few of the recording surfaces do not meet their targeted capacity, the sum of the surface capacities of all the media surfaces will be less than the target capacity, causing the entire disk drive to fail~~.

U.S. Patent Nos. 6,091,559 and 5,596,458 provide ~~for recording at~~ different BPI on different disk~~recording~~ surfaces, however these approaches~~,~~ ~~such methods~~ do not take into consideration multiple constraints, including head performance across the stroke per disk surface ~~affecting disk drive capacity~~, ~~disk drive~~ performance requirements such as ~~(e.g.,~~ throughput~~)~~ and manufacturing requirements such as ~~(e.g.,~~ test time~~)~~. Instead, disk surface z~~Z~~one frequencies are selected based on ~~measurement of~~ a single metric for~~on~~ one head.

There is, therefore, a need for ~~a method of~~ storing data in a disk drive which improves disk drive storage capacity and yield and ~~while~~ accounts for ~~meeting the desired target drive capacity or increasing the drive capacity while meeting a desired drive yield by taking advantage of the~~ head performance variation.

Summary of the Invention

The present invention satisfies thiese needs.

In an embodiment, a variable BPI storage format is a function of ~~storage~~ zones in data storage devices, such as disk drives, based on head performance variation between different heads in a set of data storage devices.

In another embodiment, ~~According to one embodiment of the present invention,~~ a population of disk drives is selected, and head performance measurements are taken for ~~each selected~~ disk~~media~~ surface locations at different frequencies. Head p~~P~~erformance distributions ~~are~~ obtained from the head performance measurements~~d data,~~ provide storage formats~~and a format optimizer uses the distributions to~~ for the disks by determining~~obtain a design~~ ~~of~~ different read/write frequencies for~~across~~ the ~~disk~~~~media surface~~ zones,~~,~~ and ~~determine head allocation. Once the different frequencies for the zones have been determined, then in each disk drive,~~ the heads in each disk drive are assigned to the ~~predetermined~~ frequencies ~~optimized for. As such, the present~~

~~invention allows maintaining consistent performance (both sequential and random throughput) across a population of disk drives, and reduced test time.~~

T~~This is accomplished by determining head performance and design of format at development/design time, and assignment of heads to different frequencies at manufacturing time. Therefore, predetermined design of formats is performed off-line, and then marries to a manufacturing test process for assignment of heads to different frequencies.~~

~~In one example, the density/format for each recording surface zone and the number of heads allocated to each density, are preselected at design time, and at manufacturing time heads are assigned to higher/lower density formats. Unlike conventional methods,~~ he head allocations and assignments are~~is~~ per head per zone, taking into consideration head performance variation across the zones. For instance~~As such~~, if- a first head ~~that~~ performs well at the inner diameter (ID) of the disk but poorly at the outer diameter (OD) of the disk, and a second head has reverse performance, then ~~that performance is traded off wherein~~ the first head is assigned a ~~to~~ high BPI~~density~~ at the ID and at low BPI~~density~~ at the OD, and the second head is assigned in the opposite fashion. T~~In~~ the per zone variable BPI storage format ~~for~~ improves~~ing~~ storage capacity by taking ~~according to the present invention,~~ several manufacturing and customer constraints ~~are taken~~ into consideration. Performance of each head across the stroke of the disk surface, as well as performance variation from one head to another, ~~is~~ determines ~~utilized in designing~~ the storage~~density~~ format and ~~assignments of~~ the head assignments~~s to the density formats~~.

In another embodiment, the head performance and the storage format are determined off-line at development/design time, and then the heads are assigned to the different frequencies at manufacturing time. For~~In one~~ example, the storage format for each ~~disk surface~~ zone and the number of heads allocated to

each data density are preselected at design time, and then the heads are assigned to high/low data density storage formats at manufacturing time.

~~The present invention provides a variable BPI storage format as a function of storage zones in storage devices, such as disk drives, based on transducer head performance variations between different heads in a set of disk drives.~~

In another embodiment, a~~The present invention provides a~~ method ~~of~~ defines~~ing~~ the~~such a~~ storage format in ~~multiple~~ data storage devices, with each data storage device having ~~a~~ multiple~~plurality of~~ storage media and ~~a plurality of~~ corresponding ~~data transducer~~ heads, each ~~transducer~~ head for recording on and playback of information from a corresponding storage media~~medium~~ in multiple zones, and ~~wherein~~ each zone including~~es~~ ~~a plurality of~~ concentric tracks for recording on and playback of information. The method includes ~~the steps of:~~ (1) selecting ~~a plurality of~~ a sample of the~~said~~ data storage devices,~~;~~ (2) for each selected data storage device, measuring a record/playback performance ~~capability~~ of each head at one or more read/write frequencies per zone,~~;~~ (3) ~~based on said performance capability measurements,~~ generating head performance~~storage density~~ distributions ~~corresponding to at least a number of the heads in said selected data storage devices~~based on the head performance measurements,~~;~~ (4) selecting a group of read/write frequencies for the~~said~~ ~~multiple~~ data storage devices, two or more frequencies for each zone, based on the~~said~~ head performance~~storage density~~ distributions,~~;~~ and thereafter, during manufacturing, (5) assigning one of the~~said~~ read/write frequencies to each head based on the performance ~~capability~~ of that head ~~per storage device~~.

Advantageously, the present invention provides consistent performance (both sequential and random throughput) across a population of disk drives, improves storage capacity and yield and reduces test time.

Brief Description of the Drawings

These and other features, aspects and advantages of the present invention will become understood with reference to the following description, appended claims and accompanying figures where:

FIG. 1A shows a~~n example~~ ~~partial schematic diagram of a~~ disk drive with a~~n example~~ data storage format~~ according to the present invention~~;

FIG. 1B shows~~ another~~ ~~example schematic of~~ drive~~the~~ ~~disk drive of FIG. 1A illustrating disk drive~~ electronics for the disk drive;

FIG. 1C shows ~~a~~ servo tracks and data tracks on a disk ~~n example~~ surface ~~format for data storage according to the present invention~~;

FIG. 1D shows a~~n example diagram representing the~~ ~~general~~ zone format~~layout~~ in~~of~~ the~~a~~ disk drive with N disks, ~~and~~ 2N heads~~,~~ and ~~depicting~~ different heads in ~~a section of~~ a zone on different disk surfaces;

FIG. 1E shows another ~~example of capacity~~ zone format~~layout~~ on a disk surface~~ a disk~~;

FIG. 1F shows ~~example of a series of radial~~ zones on a disk surface that~~,~~ ~~wherein~~ each ~~zone~~ includes ~~multiple~~ virtual cylinders;

FIG. 1G shows a~~n example representative~~ data track format~~layout~~ for the ~~each of several~~ virtual cylinders in a zone on different disk surfaces with corresponding heads;

FIG. 1H shows a ~~nother example~~ servo track and data track format~~layout~~ for a zone on different disk surfaces with corresponding heads in which~~, wherein~~ the number of servo tracks and data tracks in different virtual cylinders of a zone on different disk surfaces are the same;

FIG. 1I shows another ~~example layout wherein the~~ servo and data track format~~layout~~ that varies from zone to zone on a disk surface;

FIG. 2A shows a~~n example~~ function and flow/~~functional~~ diagram for~~of~~ ~~embodiment of steps of~~ generating the format~~layout~~ of FIG. ~~1I~~ 1A~~ according to the present invention~~;

FIG. 2B shows a graph of playback error measurement for a head at a zone at different recording frequencies;

~~FIG. 2C shows an example joint BPI distribution plot;~~

FIG. 2C~~D~~ shows a~~n example~~ histogram of the frequency capabilities of the heads in a set of disk drives at a zone at a fixed target error rate;

FIG. 2D shows a joint BPI distribution;

FIG. 3 shows a~~n example~~ flowchart ~~of an embodiment of steps~~ of vertical~~ariable~~ zoning data collection ~~process~~ in~~of~~ FIG. 2A;

FIG. 4 shows a~~n example~~ flowchart of ~~an embodiment of steps of~~ vertical zoning ~~post-measurement processing~~ and per zone joint BPI distribution extraction ~~process~~ in~~of~~ FIG. 2A;

FIG. 5 shows a~~n example~~ flowchart of head~~an embodiment of steps of vertical zone~~ assignments in ~~process of~~ FIG. 2A; and

FIG. 6 shows a~~n example~~ flowchart of ~~an embodiment of steps of~~ format generation and optimization ~~interaction process of~~ in FIG. 2A.

Detailed Description of the Invention

Data storage devices used to store data for computer systems include, for example, ~~hard~~ disk drives, floppy ~~disk~~ drives, tape drives, optical and magneto-optical drives~~,~~ and compact ~~disk~~ drives. Although the present invention is illustrated by way of a~~n exemplary magnetic hard~~ disk drive~~ 100~~, the present invention can be used in other data storage devices and other storage media~~ and drives~~, including non-magnetic storage media, is as apparent to those~~one~~ of ordinary skill in the art and without deviating from the scope of the present invention.

~~Referring to~~ FIGs. 1A-1C show~~,~~ ~~an exemplary~~ hard disk drive 100 ~~is~~ diagrammatically depicted for storing user data and/or operating instructions for a host computer ~~system~~ 54. The ~~hard~~ disk drive 100 includes~~comprises~~ an electro-mechanical head-disk assembly 10 that ~~shown in FIG. 1A as~~ includes~~ing~~ one or more rotating data storage disks 12 mounted in a stacked, spaced-apart relationship upon a ~~rotating~~ spindle 13~~. The spindle 13 is~~ rotated by a spindle motor 14 at a predetermined angular velocity.

Each disk 12 includes~~defines~~ at least one disk~~media~~ surface 23, and usually two disk~~media~~ surfaces 23 on opposing sides~~ of each disk 12~~. Each disk~~media~~ surface 23 has associated ~~is coated is coated with~~ magnetic ~~or other~~ media for recording data. The spindle ~~drive~~ motor 14 rotate~~turn~~s the spindle 13 ~~in order~~ to move the disks 12 past the magnetic transducer heads 16 suspended by the suspension arms 17 over each disk~~media~~ surface 23. Generally, each ~~magnetic~~ head 16 is attached to a~~the~~ suspension arm 17 by a head gimbal assembly (not shown) that enables the ~~magnetic~~ head 16 to swivel to conform to a~~the~~ disk~~media~~ surface 23~~s on the disks 12~~. The suspension arms 17 extend radially from a rotary voice coil motor 20~~actuator (not shown)~~. The voice coil ~~An actuator~~ motor 20~~54~~ rotates the suspension ~~actuator and head~~ arms 17 and thereby positions the ~~magnetic~~ heads 16 over the appropriate areas of the disk~~media~~ surfaces 23 in order to ~~locate and~~ read from or write ~~data from or~~ to the disk~~storage~~ surfaces 23. - Because the disks 12 rotate at relatively high speed, the ~~magnetic~~ heads 16 ride over the disk~~media~~ surfaces 23 on a cushion of air (air bearing).

–Each ~~magnetic~~ head 16 includ~~comprise~~s a read element (not shown) for reading ~~magnetic~~ data from ~~on~~ a disk ~~magnetic storage media~~ surface~~s~~ 23 and a write element (not shown) for writing data to a disk~~on the media~~ surface~~s~~ 23. Most preferably, the read element is a magneto-resistive or giant magneto-resistive sensor and ~~although not necessarily,~~ the write element is inductive and has a~~n electrical~~ write~~ing~~ width which is wider than a~~n electrical~~ read~~ing~~ width of the read element~~, which is preferably of magnetoresistive or giant magnetoresistive material~~.

E~~Referring to FIG. 1C, e~~ach disk~~media~~ surface 23 is divided into ~~a plurality of~~ concentric circular data tracks 30 that each have individually addressable data sector~~portions~~ 35~~, such as sectors,~~ in which user data is stored in the form of magnetic bits. The data sectors 35 are separated by narrow embedded ~~narrow~~ servo sectors ~~or spokes~~ 25 arranged in radially extending

servo spokes. The servo sectors 25 ~~which~~ include a series of phase-coherent digital fields followed by a series of constant frequency servo bursts. The servo bursts are radially offset and circumfer~~r~~entially sequential, and are provided in sufficient numbers ~~such~~ that fractional amplitude read signals generated ~~picked up~~ by the head 16~~read element~~ from portions of at least two servo bursts passing under the head 16~~read element~~ enable the controller 57 to determine and maintain proper ~~head~~ position of the head 16 relative to a data track 30. A~~One example of a~~ servo burst pattern for use with a head that includes ~~an inductive write element~~a /magneto-resistive read element and an inductive write element ~~head 16~~ is describe~~provide~~d by commonly assigned U.S. Patent No. 5,587,850~~,~~ entitled~~:~~ "Data Track Pattern Including Embedded Servo Sectors for Magneto-Resistive Read/Inductive Write Head Structure for a Disk Drive"~~,~~ which is incorporated herein by reference.

The ~~drive~~ controller 57 controls ~~operation of~~ the ~~pairs of magnetic~~ heads 16 ~~and media surfaces 23~~ to read from and write ~~data on~~to the ~~each~~ disk~~media~~ surfaces 23. The ~~drive~~ controller 57 preferably is~~comprises~~ an application specific integrated circuit~~s~~ chip (ASIC) ~~chip~~ which is connected by a printed circuit board 50 to ~~with~~ other ASIC~~chip~~s, such as a read/write channel ~~chip~~ 51, a motor~~s~~ driver ~~chip~~ 53~~,~~ and a cache buffer ~~chip 55, into an electronic circuit as shown in FIG. 1B~~. The controller 57 preferably includes an interface 59 which connects to the host computer 54 via a known bus ~~structure~~ 52~~,~~ such as an ATA or SCSI bus.

The controller 57 executes embedded or system software includ~~compris~~ing programming code that monitors and operates the disk drive 100~~controller system and driver 100~~. During a read or write~~data retrieval~~ operation, the host computer ~~system~~ 54 determines~~d~~ the ~~"~~address~~"~~ where the data is located i~~o~~n the disk drive 100. The address specifies the ~~, i.e., magnetic~~ head 16~~number~~, the data track 30~~,~~ and the data sector~~relevant portion(s)~~ 35 ~~of the track 30~~. This data is transferred to the ~~drive~~ controller 57 which maps the

address to the physical location in the disk drive 100, and in response to reading the servo information in the servo sectors 25, operates the voice coil ~~actuator~~ motor 20~~54~~ ~~and suspension arm 17~~ to position the~~a~~ ~~magnetic~~ head 16 over the corresponding data track 30. As the disk~~media~~ surface 23 rotates, the ~~magnetic~~ head 16 reads the servo information embedded in each servo sector~~spoke~~ 25 and also reads an address of each data sector~~portion~~ 35 in the data track 30.

During a read operation, w~~W~~hen the identified data sector~~portion~~ 35 appears under the ~~magnetic~~ head 16, the entire contents of the data sector~~portion~~ 35 containing the desired data is~~are~~ read. In reading data from the disk~~media~~ surface 23, the head 16~~read element (not shown)~~ senses a variation in an electrical current flowing through ~~a magnetoresistive sensor of~~ the read element ~~(not shown)~~ when it passes over an area of flux reversals on the disk surface 23~~ of the media~~. The flux reversals are transformed into recovered data by the read/write channel ~~chip~~ 51 in accordance with a channel algorithm such as partial response, maximum likelihood (PRML). The recovered data is then read into the cache buffer~~memory chip~~ 55 ~~of the disk drive 100 from~~ whe~~nce~~re it is transferred to the host computer ~~system~~ 54. The read/write channel 51 most preferably includes a quality monitor ~~function~~ which ~~enables~~ measures~~ment of~~ the quality of recovered data and ~~thereby~~ provides an indication of the data error rate. One channel implementation which employs channel error metrics is described in commonly assigned U.S. Patent No.~~,~~ 5,521,945 ~~to Knudson,~~ entitled~~:~~ -"Reduced Complexity EPR4 Post-Processor for Sampled Data Detection" which is~~,~~ incorporated herein by reference. The present invention~~i~~ uses the indication of recovered data error to ~~is used in order to~~ select linear bit~~data~~ density, track density and/or error correction code~~ levels, in accordance with principles of the present invention, as more fully explained hereinbelow~~.

~~Writing or storing data on the media surface 23 is the reverse of the process for reading data.~~ During a write operation, the host computer ~~system~~ 54 remembers the address~~es~~ for each file on the disk~~media~~ surface 23 and which

data sector~~portion~~s 35 are available for new data. The ~~drive~~ controller 57 operates the voice coil ~~actuator~~ motor 20~~54~~ in response to the servo information read back from the ~~embedded~~ servo sectors 25 ~~in order~~ to position the head 16~~a magnetic Head1~~, settles the head 16 into a writing position, and waits for the appropriate data sector~~portions~~ 35 to rotate under the head 16 to ~~perform the actual~~ write~~ing~~ the~~of~~ data. To write data on the disk~~media~~ surface 23, an electrical current is passed through a write coil in the inductive write element ~~(not shown)~~ of the head 16 to create a magnetic field across a magnetic gap in a pair of write poles that magnetizes the ~~magnetic storage media coating the~~ disk~~media~~ surface 23 under the head 16. When the data track 30 is full, the ~~drive~~ controller 57 moves the ~~magnetic~~ head 16 to the next available data track 30 with sufficient contiguous space for writing ~~of~~ data. If still more track capacity is required, another head 16 is used to write data to a data sector~~portion~~ 35 of another data track 30 on another disk~~media~~ surface 23.

T~~In one aspect, t~~he present invention increases the ~~data~~ storage capacity and yield of data storage devices, such as the disk drive 100, having ~~a plurality of~~ magnetic media surfaces, such as the disk surfaces 23~~, such as hard disk drive 100 including disks 12 covered with magnetic media~~.

V~~Overview of general method~~ vertical Z~~z~~oning

In every disk drive, there is a distribution associated with the head and disk surface ~~/media~~ pair performance~~ in that disk drive~~. The present invention takes advantage of that distribution to determine different linear bit density (BPI) recording frequency assignments for the heads, and optionally track allocation.

A~~ According to one embodiment of the present invention, a~~ set of disk drives is selected, and head performance measurements are taken for each selected disk~~media~~ surface location in the disk drives at different frequencies.

Empirical frequency capability histograms are extracted at a ~~given known~~ target performance metric from the measurement data. Head performance distributions ~~Probability cumulative distribution functions~~ (such as joint BPI ~~probability~~ distributions) are estimated from the histograms and fed into a format optimizer to obtain and design ~~(~~vertically zoned~~)~~ frequency format profiles ~~(i.e.,~~ across the stroke and the disk~~media~~ surface ~~zones)~~ as well as the optimal number of head allocations to the frequencies. Once the frequency format profiles and the optimal number of head allocations are ~~designed and pre-~~determined, during a test process, every head at every zone is assigned to one of the ~~multiple pre-determined~~ frequencies based on the head's performance ~~capability~~.

~~As such, the present invention allows maintaining consistent performance (sequential/random throughput) across several of disk drives, without introducing significant additional test time. This is accomplished by determining head performance and design of format at development/design time, and assignment of heads to different frequencies at manufacturing time. Therefore, the predetermined design of frequency format profiles and (optimal) number of head allocations are performed off-line while the assignment of heads to different frequencies is performed during the test process such as during manufacturing.~~

~~Unlike conventional methods, in the present invention head allocation and assignment is per head per zone, taking into consideration head performance variation across zones (i.e. across the stroke). As such, during head frequency assignment, if a first head performs well at ID but poorly at OD, and a second head has the opposite performance, the performance variation between the heads is traded off such that the first head is assigned to high density (frequency) at ID and at low density at OD, and the second head is assigned to low density at ID and high density at OD. In a method of per zone variable bits-per-inch (BPI or linear density) for improving capacity according to the present invention, several manufacturing and customer constraints are taken into consideration. And,~~

~~performance of each head across the stroke, as well performance variation from one head to another, is utilized in designing the density format and assignments of heads to the density formats.~~

~~Referring back to~~ FIG. 1A shows~~,~~ a storage format for the disk drive 100. Each disk surface 23 includes zones 60 that extend from one radius of the disk 12 to another radius of the disk 12, and the format of the zones 60 on each disk surface 23 is the same. The~~n example of density layout according to an embodiment of the present invention is shown. In one aspect, the present invention provides a~~ variable BPI storage format is ~~layout as~~ a function of the zones 60 on each disk surface 23 based ~~on e.g.~~ two data recording formats -- ~~(i.e., low~~ high data density and ~~high~~ low data density -- ~~)~~ that use~~tilize:~~ (1) head performance variation from one head 16 to the next head 16 in the disk drive 100, and (2) the performance ~~(~~variation~~)~~ of a given head 16 across the stroke of~~n~~ a disk surface 23.~~ Further, the present invention provides a method of generating said layout.~~

~~In this example, each zone comprises a group of tracks laid out in zones 60 between one radius and another radius on the disk surface 23, wherein the zone layout for the multiple disk surfaces in each disk drive 100 are the same.~~

The~~A~~ disk drive 100 includes the ~~N~~ disks 12 depicted as ~~(~~d~~D~~isks 1 to~~hrough Disk~~N, the heads 16 depicted as heads 1 to 2N, and the disk surfaces 23 depicted as disk surfaces 1 to 2N. E~~), e~~ach disk 12 includes~~having~~ two opposing disk surfaces 23, and each head 16 is associated with one of the disk surfaces 23~~a Surface1 and opposing Surface~~. For instance, head 1 is associated with disk surface 1 of disk 1, head 2 is associated with disk surface 2 of disk 1, head 3 is associated with disk surface 3 of disk 2, and head 2N is associated with disk surface 2N of disk N.

E~~2, wherein e~~ach disk surface 23 includes the zones 60 depicted as ~~has M zones (~~z~~Z~~ones_1 througoh ~~Zone~~M~~)~~ across its ~~the actuator~~ stroke, with~~, and one head per disk surface. For each disk surface,~~ z~~Z~~one_1~~1~~ ~~is~~ at the ~~in~~ ID and z~~,~~ ~~Z~~one_M ~~is~~ at the ~~in~~ OD. T~~, wherein t~~he radial boundaries on~~f~~ z~~Z~~one_1 of disk s~~S~~urface_1 of d~~D~~isk_1 are the same as the radial boundaries of z~~Z~~one_1 on~~f~~ disk s~~S~~urface_2 of d~~D~~isk_1, and so on. Similarly, the radial boundaries of z~~Z~~one_M on disk s~~S~~urface_1 of d~~D~~isk_1 are the same as the radial boundaries of z~~Z~~one_M on disk s~~S~~urface_2 of d~~D~~isk 1, and so on. However, different zones 60 across the stroke on each disk surface 23 need not necessarily have the same number of data tracks 30 or TPI ~~(e.g., Zone1 and ZoneM on the same surface do not necessarily have the same number of tracks)~~. For example, z~~Z~~one_1 on disk s~~S~~urface_1 of d~~D~~isk_1 ~~(i.e., Head1)~~ has the same number of data tracks 30 and the same radial~~physical zone~~ boundaries as z~~Z~~one_1 on disk s~~S~~urface_1 of d~~D~~isk N ~~(Head 2xN)~~, ~~etc.~~ a~~A~~nd~~,~~ z~~Z~~one_M on disk s~~S~~urface_1 of d~~D~~isk_1 ~~(Head1)~~ has the same number of data tracks 30 and the same radial~~physical zone~~ boundar~~y~~ies as z~~Z~~one_M on disk s~~S~~urface_1 of d~~D~~isk_N ~~(Head2N)~~. However, the number of data tracks 30 in z~~Z~~ones_1 and ~~Zone~~M can be different.

Each disk surface 23 also includes virtual cylinders 39 depicted as virtual cylinders 1 to n. Each zone 60 includes multiple virtual cylinders 39, and each virtual cylinder 39 includes multiple data tracks 30 on each disk surface 23. ~~The physical zone boundaries vertically align on the disks in each disk drive, forming virtual cylinders 39 (VC). In this example, there are n virtual cylinders 39, VC1 through VCn.~~ Further, w~~W~~ithin a virtual cylinder 39, different heads 16 may read and ~~/~~write at different frequencies (~~e.g.,~~ variable BPI) to provide~~, hence the concept of~~ vertical zoning.

~~The level of track density (TPI) can be one of fixed number of preselected levels or can be derived from an algorithm that is based on the location of a portion 35 of the media surface 23. Embedded servo sectors 25 are initially written on a media surface 23 during a factory servo-writing process at a servo~~

~~track density that can be higher than the data track density, as illustrated in FIG. 1C. Servo bursts within each servo sector 25 are provided in such number and placement to enable accurate positioning of the magnetic head 16 in a full range of positions across the media surface 23, given the particular effective width and characteristics of the read element of a particular head (the read element width typically being narrower then the writer carry out the head positioning method, information in the embedded servo sector 25 is read by the magnetic head 16 and passed to the drive controller 57 which directs the actuator motor 20 to readjust the position the suspension arm 16. In the example shown in~~ FIG. 1C shows~~,~~ the data tracks 30 and the servo tracks 37 on the disk surface 23. The data tracks 30 include the data sectors 35, and the servo tracks 37 include the servo sectors 25. ~~the servo track density is about 150% of the maximum possible data track density.~~ F~~In FIG. 1c~~five servo tracks 37 depicted as servo tracks ~~(e.g.,~~Sa, Sb, Sc, Sd and Se~~)~~ are shown in relation to three data tracks 30 depicted as data tracks Tk1, Tk2 and Tk3.

The servo tracks 37 are written on the disk surface 23 during manufacturing at a servo track density that is about 150% of the maximum data track density. -The s~~S~~ervo track density is determined by ~~determining~~ the maximum~~.~~ read width and the minimum~~.~~ write width of a population of the~~magnetic~~ heads 16. After writing the servo tracks 37~~wedges 25~~ at the servo track pitch, the ~~actual~~ data tracks 30 can be written at any ~~disk~~ radial position between the servo tracks 37. The data track density (TPI) can be selected from predetermined levels or can be based on the location of a data sector 35. Additional tests~~,~~ can be performed to determine the optimum data track density of the disk~~media~~ surface 23. Each servo track 37 comprises radially similarly situated ~~servo information in~~ servo sector~~wedge~~s 25 in the servo spokes. For example, ~~(e.g.,~~ the ~~set of~~ servo track~~information~~ Se contains servo sectors 25 at essentially the same radial distance from the ~~disk~~ center of the disk 12, the servo track Sd contains servo sectors 25 at essentially the same radial distance from the center of the disk 12, ~~form a servo track circumferentially, set of servo~~

~~information So at essentially same radial distance from the disk center form another servo track circumferrentially,~~ etc.~~)~~

FIGs. 1D to 1I show vertical zone formats in which different heads 16 on different disk surfaces 23 may read/write at different linear frequencies (variable BPI) on the data tracks 30 within a virtual cylinder 39.

FIG. 1D shows a~~n example diagram representing the~~ ~~general~~ zone~~-~~ 60 format~~layout~~ of the ~~a~~ disk drive 100~~,~~ with N disks 12,~~and~~ 2N heads 16~~,~~ and ~~depicting~~ different heads 16 in ~~a section~~ ~~z~~Zone 1 on different disks 12. FIG. 1E shows another ~~example of capacity~~ zone 60 format~~layout~~ on the disk surface 23~~a surface a disk~~. FIG. 1F shows each ~~example of a series of radial~~ zones~~,~~ 60 on the disk surface 23 ~~, Zone1 through ZoneM, on a disk surface, wherein each zone~~ includes multiple virtual cylinders 39. Zone 1 includes virtual cylinders 1 to j, and zone M includes virtual cylinders 1 to i. The radial boundaries of the zones 60 are shown as dark circles, and the radial boundaries of the virtual cylinders 39 are shown as light circles. FIG. 1G shows a~~n example representative~~ data track 30 format~~layout~~ for the ~~each of several~~ virtual cylinders 39 in a zone 60~~(e.g., Zone1),~~ on different disk surfaces 23 with corresponding heads 16. -FIG. 1H shows a data track 30 and ~~nother example~~ servo track 37 format~~and data track 35 layout~~ for a zone 60 on different disk surfaces 23 with corresponding heads 16 in which ~~, wherein~~ the number of data tracks 30 and servo tracks 37 ~~and data tracks~~ in different virtual cylinders 39 of a zone 60~~(e.g., Zone1)~~ on different disk surfaces 23 is~~are~~ the same. ~~And,~~ FIG. 1I shows a data track 30 and servo track 37 format for zones 60 on the disk surface 23 with a corresponding head 16 in which ~~another example layout wherein the servo and~~ the data track 30 and servo track 37 format~~layout~~ varies from ~~zone~~ zone 1 to zone M~~60 on a disk surface 23~~. ~~In all of the above examples of vertical zoning layout according to the present invention, within a virtual cylinder 39, different heads on different surfaces may read/write at different linear frequencies on the data tracks (e.g., variable BPI).~~

~~Overview of~~ Format Optimization

The ~~In one embodiment, a~~ vertical zoning ~~method according to the present invention~~ includes designing, ~~/~~optimizing and ~~(~~selecting~~)~~ ~~for~~ two or more recording frequency profiles ~~(i.e.,~~ per zone~~)~~ for a sample number of disk drives ~~(e.g., performed~~ off-line during the disk drive development/design phase~~)~~. Then, for a population of disk drives, in each disk drive, each head is assigned to one of the predetermined frequencies for a given zone ~~(e.g.,~~ during the disk drive manufacturing phase~~)~~. A~~The assignment step includes assigning a~~ predetermined read/write frequency (BPI) is assigned to each head based on a known number of head allocations and the head's performance~~ capability~~. A head assigned to a high~~er~~ frequency~~/density (HD)~~ records more bits on a track, and a head assigned to a low~~er~~ frequency~~/density (LD)~~ records less bits on a track.

Performance testing of the head and disk surface pairs occurs after full read/write and servo calibration and optimization of the disk drive. I~~Referring to the example in FIG. 1A,~~ if the tested performance of h~~H~~ead_1 at z~~Z~~one_1 on disk s~~S~~urface_1 of ~~-~~d~~D~~isk_1 at a given frequency ~~(after full drive read/write and servo calibration/optimization)~~ is better than a ~~desired~~ target performance metric, then ~~that (strong) head,~~ h~~H~~ead_1~~,~~ is considered strong since it is capable of ~~to have some margin for~~ storing more information than ~~it was~~ originally accounted for. Thus, the ~~designed~~ recording frequency can be increased at z~~Z~~one_1 on disk s~~S~~urface_1 of d~~D~~isk_1 for h~~H~~ead_1 yet ~~so as to ensure~~ the~~its~~ performance does not fall below the ~~desired~~ target performance metric. If the tested performance of h~~H~~ead_2 at z~~Z~~one_1 on disk s~~S~~urface_2 of d~~D~~isk_1 at the~~at~~ same frequency ~~(after full drive read/write and servo calibration/optimization),~~ is worse than a ~~desired~~ target performance metric, then ~~that (weak) head, H~~head_2 is considered weak~~,~~ but can be compensated for by relaxing the frequency at which h~~H~~ead_2 operates~~, so as~~ to ensure the target performance metric is met. Performing the above trade-off between the heads for all the zones~~, without loss of overall capacity,~~ provides ~~resulting~~ frequency profiles ~~(i.e.~~ across the stroke~~)~~ that are

~~comprising~~ vertically zoned frequency format profiles without loss of ~~overall~~ storage capacity.

Advantageously, ~~As the above example shows,~~ by compensating for h~~weak H~~ead 2, rather than failing the disk drive due to ~~the weak H~~head 2, the vertical zoning improves ~~the disk drive~~ yield. ~~Without application of vertical zoning such a disk drive would not have passed the test limits, and hence would have failed.~~ Furthermore, the format optimizer uses the head ~~estimates of~~ performance (~~i.e.,~~ read/write frequency capability) ~~cumulative~~ distributions ~~function~~ at every zone and a target performance metric~~,~~ to design a group of read/write frequency format profiles for strong and weak ~~and strong~~ heads within a given disk drive. The format optimizer also determines the optimal number of~~,~~ ~~for example,~~ strong versus weak ~~versus strong~~ heads.

The format optimizer does not determine which specific head is ~~actually~~ at the high or lowe~~r~~ ~~or higher~~ frequency, but does ~~only~~ provide~~s~~ a breakdown of the number of heads at the high~~lower~~ frequency and the number of heads at the low~~higher~~ frequency. The~~at~~ breakdown is fixed, performed off-line, and ~~is~~ used during the head assignment~~ process~~. Then, in the head assignment~~ process~~ ~~(e.g.,~~ during a manufacturing test~~ process)~~, out of 2~~x~~N heads in a disk drive with N disks, the number of heads ~~that have to be~~ assigned to each predetermined frequency/~~format~~ is ~~also pre~~determined~~ (e.g., number heads to assign to low frequency and number of heads to assign to high frequency)~~.

T~~As such, t~~he heads within a set of disk drives are allocated to the predetermined group of read/write frequencies as part of the optimization process to meet the storage capacity and yield requirements for the disk drives. The allocation process allocates a number of the heads in a disk drive to ~~one of~~ the predetermined/~~designed~~ frequencies, however the specific assignment of a particular head to a particular frequency is performed later during ~~as part of~~ the assignment process~~ thereafter~~. For ~~In one~~ example, in a two~~2~~ read/write

frequency design (high frequency ~~density~~ and low frequency~~density~~) for a set of disk drives- each with eight~~8~~ heads, in each disk drive for z~~Z~~one 1 on all the disk surfaces, any five~~3~~ ~~heads~~ of the eight~~8~~ heads are allocated to the high~~lower~~ frequency and any three~~5~~ ~~heads~~ of the eight~~8~~ heads are allocated to the low~~higher~~ frequency ~~in the allocation process~~ based on the performance measurements of ~~all~~ the heads in the set- of the disk drives. Thereafter, the specific assignment of each particular head to a particular predetermined frequency is performed~~ as part of the assignment process~~. For example,- in a first disk drive ~~heads 1, 3, 4 are assigned to the low frequency, and~~ heads 2, 5, 6, 7, 8 are assigned to the high frequency and heads 1, 3, 4 are assigned to the low frequency, whereas in a second disk drive ~~heads 2, 3, 8 are assigned to the low frequency, and~~ heads 1, 4, 5, 6, 7 are assigned to the high frequency and heads 2, 3, 8 are assigned to the low frequency. T~~, and so on, wherein t~~he specific head assignments depend on the specific capability of the heads in each disk drive.

The optimal number of heads per frequency ~~(i.e., head allocation)~~ is determined at the same time that the group of read/write frequencies are ~~designed/~~selected by the format optimizer~~,~~ by solving a joint constrained optimization problem. For example, in a ~~the 8-head~~ disk drive with eight~~8~~ heads andwith~~above, for the case of two frequencies~~ a ~~(~~high frequency ~~(freq1)~~ and a low frequency~~ (freq2)~~) that are each a different ratio of a reference frequency~~ at a ratio to a frequency freq~~, in each vertical zone, allocating two~~2~~ heads to the high frequency~~1~~ and six~~allocating 6~~ heads to the low frequency~~2,~~ provides a specific ~~first~~ storage ~~disk drive~~ capacity. Changing the~~said~~ frequency ratio~~ of the frequencies~~ and the number of heads allocated to each frequency~~,~~ provides a different storage capacity~~ for that disk drive~~. Thus, ~~As such,~~ the disk drive storage capacity is a function of the number of heads multiplied by the frequency allocated to each head per zone. —For example, if a nominal disk surface data storage ~~level capacity~~ is 1 unit, and if the high frequency ~~1~~= 4/3 x the reference frequency and the low frequency~~2~~ = 2/3 x the reference frequency, then one

head can be at the high frequency~~1~~ for every one head at the low frequency~~2,~~ to maintain ~~whereby~~ the average disk surface data storage ~~capacity~~ at~~is~~ 1 unit.

The head performance distributions represent percentages of the heads in the disk drives than can operate at different frequencies. For example, the head performance distribution i~~a~~s a BPI distribution that represents the head frequency capability at a target performance metric. –Using the head performance distributions (~~i.e.,~~ the head read/write frequency capability distributions at the target performance metric for every zone), the number of heads, the format of the virtual cylinders~~,~~ and the~~a~~ desired storage capacity, the format optimizer determines the frequency for each virtual cylinder in each zone ~~in each virtual cylinder~~ and the number of heads in each disk drive allocated to each frequency~~,~~ ~~in order~~ to achieve the~~at~~ desired storage capacity. Thereafter, in the assignment process ~~(e.g.,~~ as part of ~~a~~ testing ~~of~~ each disk drive~~)~~, each ~~specific~~ head in a population of disk drives is assigned to one of the predetermined frequencies based on the allocation criteria and the specific head performance. For ~~an~~ example, in a ~~4-head~~ disk drive with four~~4~~ heads, the format optimizer considers three~~3~~ heads at the high frequency~~1~~ and one~~1~~ head at the low frequency~~2~~, then two~~2~~ heads at the high frequency~~1~~ and two~~2~~ heads at the low frequency~~2~~, and then one~~1~~ head at the high frequency~~1~~ and three~~3~~ heads at the low frequency~~2~~. Thus, ~~And in each case,~~ the format optimizer uses~~ing~~ the ~~estimated~~ head performance ~~cumulative~~ distribution~~functions~~ to determine~~s~~ the ~~disk drive yield. The head performance distributions (e.g., BPI distributions) represent percentages of the heads in the disk drives than can operate at different frequency densities. This allows the format optimizer to determine the~~ storage capacity and yield~~ and capacity.In the description of the example embodiment herein, head performance distribution, such as a BPI distribution, represents head frequency capability (probability) cumulative distribution at a target performance metric.~~

In one version of the optimization process, the ~~disk drive~~ yield is maximized while meeting a constraint on storage capacity. In another version, the storage capacity is maximized while meeting a constraint on ~~disk drive~~ yield. In the former case, the format optimizer uses a format where the maximum number of disk drives qualify and the fewest number of disk drives fail to reach the required storage capacity~~higher frequency is freq1 and the lower frequency is freq2~~. For ~~In the~~ example, in a ~~4-head~~ disk drive with four4 heads and a nominal disk surface data storage~~capacity~~ of 1 unit, allocating two~~2~~ heads to the high frequency~~1~~ and two~~2~~ heads to the low frequency~~2,~~ provides a nominal data ~~surface capacity of 1 unit, and~~ storage ~~capacity~~ of 4 units~~for the disk drive~~. In the later case, the format optimizer uses a format where the maximum number of disk drives reach the required storage capacity and the fewest number of disk drives fail to qualify. For example, in a disk drive with four4 heads and a nominal disk surface data storage~~capacity~~ of 1 unit, allocating three~~If 3~~ heads ~~are allocated~~ to the high ~~higher~~ frequency~~1~~ and one~~1~~ head to the low ~~the lower frequency~~ frequency~~2~~ provides~~,~~ a higher data storage ~~capacity~~ of ~~(i.e. 4~~.66~~ and 2/3)~~ units~~is achieved for the disk drive. In that case, the format optimizer lowers the zone recording frequencies to meet that constraint of capacity of 1 unit per surface. As such, for 2 heads at freq1 and 2 heads at freq2, the format optimizer manipulates the difference of those frequencies such that surface capacity always reaches 1 unit, but maximizes yield whereby a maximum number of disk drives qualify and fewest number of disk drives fail to reach required capacity.~~

Thus, ~~As such, according to one embodiment of the present invention, a~~ the vertical zoning ~~approach~~ for variable BPI ~~design~~ includes ~~use of~~ an off-line predetermined per zone format design ~~of formats~~ based on disk drive data collection and head performance ~~(joint) BPI~~ distribution extraction~~methods~~. -In one version, a fixed predetermined zone boundary format~~layout~~ is used to design multiple frequency BPI formats based on representative or actual joint BPI distributions at one or more desired target performance metrics (such as off-track

symbol error rate) and~~, wherein~~ the joint BPI distributions are extracted from a finite pre-selected set of disk drives.

The collected data is used to extract the~~said~~ joint BPI distributions for the heads at every ~~(pre-selected)~~ zone, and the per zone design of ~~low/~~high and low data density formats for the heads is performed off-line. The format optimizer solves a constrained joint optimization~~ process~~ off-line to obtain the~~said~~ format designs~~,~~ using well-known constrained optimization routines. Using joint BPI distributions allows consideration of potential correlation of BPI capability of the heads across the stroke as well as the individual contribution of each head to the storage ~~overall drive~~ capacity ~~(or areal density)~~ and yield.

The off-line format design ~~of formats~~ allows the format optimizer to consider~~ation of~~ ~~other potential constraints that may arise, as~~ additional constraints~~ within the optimizer, and hence solved by the optimizer~~. For example, as more information is obtained in quantifying the thermal stability constraints of the disks~~recording media~~ (which in turn places an upper bound on~~of~~ linear bit density for the heads) the off-line format design does ~~provides the ability of~~ not exceed~~ing~~ the~~o~~se constrain~~limi~~ts. Likewise, i~~I~~f there are data rate constraint~~limitation~~s in either the write process ~~capability~~ or the ASICs ~~component capability~~, such constraints may be cast within the joint constrained format optimizer to ensure the~~said~~ constrain~~limi~~ts are not exceeded.

Data ~~Overview of~~ Measurement~~ Process~~

A~~In one implementation of the method of the present invention, a~~ measurement procedure is used to collect data~~,~~ from which ~~such that after processing of the collected data,~~ one-dimensional (1-D), two-dimensional (2-D) and ~~as well as~~ three-dimensional (3-D) ~~joint~~ BPI ~~(probability)~~ distributions at a desired read/write target error rate (or any other ~~choice of~~ metric) can be extracted. -Data is collected based on head capability measurements taken at different radial positions on~~of~~ the disk. ~~The collected data is used to extract 1-D,~~

~~2-D and 3-D empirical distributions at a target choice of performance metric.~~ The ~~dimensions are dimensions of the distributions, and the~~ distributions represent the capability of each head at different radial positions. For example, several disk drives which collectively include 1000 heads are selected for measurement, and ~~. In a measurement process,~~ record/playback error rate measurements of the 1000 heads from z~~Z~~one 1 to z~~Z~~one 24 of the disk surfaces at different frequencies are obtained. Thereafter, in post-measurement data processing~~:~~ (a) the BPI capability of each~~very~~ head at a fixed target performance metric at ~~e.g.~~ z~~Z~~one 1 is determined ~~in order~~ to obtain a 1-D BPI distribution, (b) the BPI capability of each~~the~~ head at a fixed target performance metric at ~~e.g.~~ z~~Z~~ones 1 and ~~zZone~~ 5 is determined ~~in order~~ to obtain a 2-D BPI ~~joint~~ distribution, and (c) the BPI capability of each~~the~~ head at a fixed target performance metric at ~~e.g.~~ z~~Z~~ones 1, ~~zZone~~ 5 and ~~zZone~~ 20 is determined ~~in order~~ to obtain a 3-D BPI ~~joint~~ distribution.

The BPI distributions are then passed~~used as input~~ to the format optimizer to solve three constrained optimization problems to provide head frequency per zone allocations. The three constrained optimization problems~~, wherein:~~ (1) ~~one problem~~ maximize~~s~~ the ~~disk drive~~ yield while preserving the ~~same~~ storage~~drive~~ capacity, (2) ~~another~~ maximize~~s~~ the storage~~disk drive~~ capacity while preserving the ~~same drive~~ yield~~-and~~, and (3) ~~another~~ maximize~~s~~ the ~~disk drive~~ yield while ensuring a ~~desired~~ target storage ~~drive~~ capacity is met at a fixed target ~~track-per-inch (~~TPI~~)~~. C~~Additionally, c~~ustomer related or ~~application specific integrated circuit (~~ASIC~~)~~ data rate ~~(limitation)~~ constraints are also consider~~utiliz~~ed. -The format optimizer can ~~is capable of~~ solve~~ing~~ any ~~one~~ of these ~~above-mentioned~~ three problems, and ~~wherein~~ one problem can take priority over another depending on the process phase. For example, at an earliy~~er~~ development phase ~~of a program~~ where the disk drive components are not matured~~-yet~~, meeting the storage ~~drive~~ capacity may be a challenge. In that phase~~case~~, the format optimizer can ~~be used to~~ design the variable BPI format profiles by solving the second problem~~-above~~. Then, a~~A~~s the disk drive components mature, ~~and~~

~~such that~~ meeting the storage~~drive~~ capacity becomes easier and meeting the ~~drive~~ yield becomes more important, the first problem may be solved ~~considered instead~~. Thereafter, as part of a test process, an assignment algorithm ~~is used to~~ ensures the appropriate head assignments to the predetermined high and ~~/~~low data density ~~(pre-specified)~~ formats per head and per zone or across the ~~head~~ strokes~~,~~ based on the head allocation breakdown of the format optimizer.

The y~~Disk drive y~~ield is improved while meeting the ~~desired~~ target storage ~~drive~~ capacity by allowing a frequency format ~~layout~~ with ~~(e.g.,~~ high and low frequencies and ~~y) with~~ a predetermined number of high and low performing head allocations. Us~~tiliz~~ing realistic constraints such as ASIC data rate limitations, the same fixed target TPI is maintained by increasing the average target BPI across the ~~head~~ stroke ~~on a disk surface~~ to achieve the target~~desired disk drive data~~ storage capacity. As such, head performance variation from one head to the next head in the disk drive ~~(~~and for ~~across~~ the head ~~stroke~~ across the stroke of the~~a~~ disk surface~~)~~ is us~~tiliz~~ed to ~~allow~~ increase~~ing~~ the storage capacity ~~areal density of the stored information~~ while preserving the ~~same overall disk drive~~ yield. For ~~In one~~ example, the~~a~~ vertical zoning format~~layout~~ use~~method according to the present invention utilize~~s several design constraints to improve ~~the drive~~ yield using a variable high and low~~/high~~ BPI design with a fixed predetermined number of head allocations as a function of the zones while meeting the target storage capacity at a fixed target TPI. The h~~H~~ead performance variation or correlation across the stroke is also us~~tiliz~~ed.

Further, ~~the method of present invention takes into consideration~~ the difference in data storage~~storage capacity~~ of two or mores zones on a disk surface is considered~~,~~ as it affects storage ~~overall disk drive~~ capacity. The storage ~~disk drive~~ capacity is defined as a weighted combination of the zone capacities across the stroke on each disk surface in the disk drive. A correlation in the head performance statistics is extracted from one head to another head,

and for every ~~given~~ head considered in a set of disk~~s~~ drives across the ~~head~~ stroke on each disk surface.

The joint constrained optimization ~~process~~ determines a per zone target high and low ~~low/high~~ data density format~~/layout~~. The optimization ~~process~~ takes into account constraints including customer related requirements such as ~~the requirement of a~~ minimum logical block count ~~(LBA)~~, monotonic data rate, and maximum data rate ~~requirements~~ at the outer zone~~-area~~s which can be formulated into ~~(~~additional~~)~~ constraints. ~~Head allocation and assignment according to the present invention improves manufacturing yield and provides a disk drive with minimal performance degradation (i.e., sequential or random throughput as well as test process time).~~

Example Implementation

FIG. 2A shows a~~n example~~ function and flow diagram ~~of an example implementation of the above-described method according to the present invention~~ for generating the optimal data density format~~/layout~~ shown ~~by example~~ in FIG. 1A. -The function and flow diagram~~example method in FIG. 2A~~ includes~~:~~ a data ~~collection/~~measurer~~ment~~ ~~process (block)~~ 62, a post-measurement data ~~post-measurement~~ processor ~~(block)~~ 64, a format optimizer ~~process (block)~~ 66 and a~~,~~ format generator ~~process and head assignment process (block)~~ 68~~,~~ ~~example embodiments of which are described below~~.

Data ~~Collection/~~Measurer~~ment process~~

The data measurer 62 takes data measurements for every zone at a finite number of frequency samples.

T~~In one embodiment, t~~he data measurer~~ment block~~ 62 implements a measurement procedure that includes~~ing~~ the steps of:

(1)~~.~~ Create~~ing~~ several different predetermined linear bit density format profiles includ~~compris~~ing a profile of different frequencies per zone across the

~~actuator~~ stroke, such as ~~e.g.~~ a first profile including high frequency ~~freq1~~1 for z~~Z~~one_1, high frequency ~~freq1~~2 for z~~to Z~~one 2~~,~~ ..._.~~,~~ high frequency ~~1~~M for z~~Z~~one_M~~;~~ and a second profile including low frequency~~2~~_1 for z~~Z~~one_1, low frequency~~2~~_2 for- z~~Z~~one 2~~,~~ ..._.~~,~~ low frequency~~2~~_M for z~~Z~~one_M~~, etc.,~~ to be loaded on a representative number of disk drives selected for the measurement process (or if possible on all ~~of~~ the available disk ~~built~~ drives for that build);

(2)~~.~~ __-Loadi~~ng~~ a frequency format profile;

(3)~~.~~ __Performi~~ng~~ read/write and servo optimization and calibration;

(4)~~.~~ __Tak~~e~~i~~ng~~ head performance measurements including ~~e.g. (~~off-track~~)~~ mean square error ~~(MSE)~~ or quality metric ~~(QM)~~ and/or symbol error rate ~~(SER) measurements~~ at pre-selected frequencies for preferably all available zones~~,~~ and sav~~e~~i~~ng~~ the data; and

(5)~~.~~ __-Repeati~~ng~~ steps 2-4 ~~above~~ for ~~all~~ the remaining frequency format profiles.

The above steps are performed for the selected disk drives in the measurement process.

Thus, in the disk drive 100,~~As such density is selected and~~ the data is recorded on a data sector 35~~portion~~ of the disk~~media~~ surface 23 at the selected data density by positioning the~~a magnetic~~ head 16 abutting the data sector ~~portion~~ 35 ~~of the diskmedia surface 23,~~ and sending the appropriate write signals to the ~~write element (not shown) of the magnetic~~ head 16. Typically, a sample of data is recorded on the disk surface 23 such that a significant number of errors are detected (such as ~~e.g.,~~ ten ~~10~~ errors per error rate measurement)~~, is recorded on the diskmedia surface 23~~ to obtain a statistically representative sampling of the error rate for the data sector ~~portion~~ 35 ~~of the diskmedia surface 23~~. _Thereafter, the recorded data is read by the ~~read element (not shown) of the magnetic~~ head 16~~,~~ and ~~the data read is~~ stored by the host computer ~~system~~ 54 for evaluation. An error rate of the recorded data is measured or compiled by comparing the ~~actual~~ written data with the read data, element-by-element. The

~~Suitable methods of determining the~~ error rate can be determined using a ~~include actual~~ bit error measurement in which a bit of data read from the disk~~media~~ surface 23 is compared with the correct bit, a~~or~~ bit steam measurement in which a bit stream of data read from the disk surface 23 ~~a correct bit stream in~~ is compared with a correct bit stream, or a ~~measured bit stream. An alternative method uses the~~ mean square error metric measurement ~~method~~ in which a waveform read from the disk~~media~~ surface 23 is compared with an ideal waveform to provide an error signal that is squared and summed to form the error metric.

In this description, a component distribution is defined as~~to be~~ a ~~(~~random~~)~~ variation (~~i.e.,~~ tolerance) of a pre-specified ~~(~~target~~)~~ nominal component parameter such as a head write/read width, and a ~~the term~~ distribution is defined as a~~the~~ probability distribution function~~ (PDF)~~. During the early product development~~ process~~, when the head performance distributions are wide and unreliable, data from a mature~~d~~ set of disk drives is used for extracting reference ~~(~~joint~~)~~ BPI distributions at a target performance metric such as on-track symbol error rate, off-track symbol error rate, on-track mean square error ~~(e.g., off/on track error rate~~ or off-track mean square error~~)~~. Later~~ in the process~~, when the ~~amount of~~ head performance variation~~s~~ from one phase to the next in the distribution is expected to be minimal, new sets of measurement data are collected using a selected ~~plurality/~~population of disk drives at their more mature~~d~~ stages.

Thus, a number of BPI formats including the nominal target format are selected. Then, on-track or off~~- or on~~-track symbol error rate or mean square error ~~MSE~~ measurements are taken at different pre-selected locations of the disk surfaces, such as the ~~e.g.,~~ outer, middle and inner zones. The performance measurements can be limited to ~~-are taken (in one example scenario described further below, the choice is limited to~~ these~~said~~ three zones~~,~~ to reduce the measurement time~~ for performing measurement)~~. However, preferably the

performance measurements over multiple zones ~~can be performed~~ and other measurements such as off-track ~~measurements (e.g.,~~ 747 ~~measurements)~~ can be performed~~also be taken~~. The nominal formats are generated from the data.

Two or more different linear bit density format profiles can be loaded at a time. In one example, two variable BPI format per zone design (~~low/~~high and low data density format profiles) can be created for ~~the purpose of~~ measurement data collection during every build. In this way, more statistical data can be collected from more disk drives, however~~,~~ there will be only two frequency samples per zone available for ~~data~~ post-measurement data processing.

P~~Raw Data~~ post-M~~m~~easurement Data P~~-p~~rocessor

P~~In the above steps, measurements (e.g. either MSE or SER) for every zone are taken at a finite number of frequency samples. In~~ post-measurement data processor ~~processing (post-processing) block~~ 64~~,~~ uses~~ing~~ the available performance metric~~,~~ measurements ~~are used~~ to calculate each head's frequency performance, for instance as ~~(e.g.,~~ kilo flux per inch (kFCI) or kilo bits per inch (kBPI),~~)~~ at a ~~given~~ target performance metric. The performance of every head at every zone is determined as a function of the~~said~~ read/write frequency profiles used for the measurements.

For example, if six~~6~~ different frequency profiles are used, then for every head per zone, the data measurer~~ment process~~ 62 provides measured data as a function of six~~6~~ frequency samples at a target performance metric. In the post-measurement data processor 64, ~~all~~ the measured data is sorted and the performance of every head at every zone at the six~~6 different~~ frequency samples is extracted to generate frequency capability histograms at a target performance metric~~(e.g. error rate)~~.

~~Referring to~~ FIG. 2B shows a graph of playback error measurement for a head at a zone at different recording frequencies. The curve shows head

performance as a function of frequency (BPI). T~~, the samples can be depicted in a two dimensional graph wherein the~~ x-axis ~~(horizontal)~~ is the read/write frequency in ~~(e.g., frequencies 1 through 6,~~ kBPI at the outer diameter~~OD)~~, and the y-axis ~~(vertical)~~ is the on-track symbol error rate o~~measurement~~ in a log scale ~~for Head1 at zone 1 for each of~~ the~~said 6 frequencies.~~ -E~~(log SER) (in FIG. 2B, each~~ frequency~~sample data~~ sample 70 is depicted as ~~a~~ "+", each curve fit point 72 is depicted as "o" and each projected frequency 74 is depicted as "◊"~~)~~.

In the illustration, head 1 at zone 1 in disk drive 3 is measured at six frequency samples. The curve is generated using a least square polynomial fit to the six frequency samples. The projected frequency (BPI) for a target on-track symbol error rate is extracted from the curve by interpolation or extrapolation. For example, if the target on-track symbol error rate is 10^{-8} then the projected frequency is determined by interpolation, whereas if the target on-track symbol error rate is 10^{-6} then the projected frequency is determined by extrapolation. The on-track symbol error rate varies ~~measurement can vary e.g. from 1e^{-4} (i.e., 1x10^{-4}) to 1e^{-7} (i.e., 1x10^{-7})~~ as a function of ~~the 6~~ frequency~~ies,~~ and ~~wherein the error rate~~ increases as the ~~read/write~~ frequency ~~(density)~~ increases.

The nominal kBPI (before vertical zoning) and the kBPI gain relative to the nominal kBPI are also shown. Head 1

can be classified as a strong head because there is reasonably significant margin before its on-track symbol error rate of –9.1 (log) at a nominal frequency/kBPI of ~ 188 can be changed to a projected on-track symbol error rate of –6.22 at a frequency/kBPI of ~ 217. Hence, there is a total kBPI gain of ~ 29, allowing the nominal frequency to increase by 15% while meeting the target on-track symbol error rate performance metric of $6x10^{-7}$. Thus, head 1 of disk drive 3 has a frequency capability of about 217 at the target on-track symbol error rate of $6x10^{-7}$, which provides a sample for the generation of a histogram.

FIG. 2C shows a histogram of the frequency capabilities of the heads in a set of disk drives at a zone at a target performance metric. The histogram 76 is constructed using the projected frequencies determined in FIG. 2B for the heads in the selected disk drives reading from zone 1 at the target on-track symbol error rate of $6x10^{-7}$. The x-axis is the projected frequency capability at the outer diameter, and the y-axis is the number of heads. The histogram is extracted and empirical, has a normal distribution fit and has a width that corresponds to the head performance variation.

Additional histograms ~~To determine frequency capability of e.g. Head1 at Zone1 at a target error rate of $1e^{-6}$ (i.e., $1x10^{-6}$), a curve is fit (e.g., using known curve-fitting techniques such as least squares polynomial fit) to the 6 samples, to determine by interpolation the frequency value that gives rise to that target error rate (in FIG. 2B, each curve fit point 72 is depicted by a "o"). If the target error rate is at $1e^{-8}$ (i.e., $1x10^{-8}$) then the frequency value that gives rise to that target error rate is determined by extrapolation (in FIG. 2b, the projected or extrapolated frequency value 74 is shown as a diamond shape). The process for that target error rate is performed for Zone1 for all the heads in the selected disk drives used in the measurement process 62, to create a histogram (FIG. 2D) of the frequency capabilities of all the heads in the disk drives at Zone1 at that fixed target error rate.~~ are constructed ~~The process is the same~~ for ~~all~~ the remaining zones based on the frequency capabilities determined from the graphs based on the performance measurements taken at the remaining zones so that ~~.~~ ~~As such, using~~ every available head considered in the disk drives under measurement has~~,~~ BPI histograms ~~can be extracted~~ at a ~~given~~ target performance metric per zone.

~~FIG. 2B shows an example curve of error rate of performance (SER) as a function of BPI (e.g., SER at 6 different BPI/frequency samples) for a head located at the outer diameter (OD) of the disk. Also shown is extracted BPI/frequency capability value of that head at a zone (e.g., OD) for the specified~~

~~target error rate, using interpolation/extrapolation (i.e., if the specified desired target error rate is outside the performance range, extrapolation or interpolation, such as polynomial fit, is used as necessary). The amount of BPI gain, or margin relative to the nominal BPI setting, is also specified and marked.~~

Thus, ~~The~~ performance measurements are provided for each head at each zone in the selected disk drives, the graphs are generated for each head at each zone, the frequency capabilities for each head at each zone are determined for a target performance metric, and the histograms are constructed for each head at each zone for the target performance metric. Likewise, ~~above process is performed for all the heads considered in the measurement procedure, and a BPI/frequency capability of all heads at a given target error rate for every zone is generated. Thus, in this manner, BPI/frequency capability histograms for every zone at a specified target error rate are constructed.~~ I~~i~~f a~~the~~ histogram of head BPI capability at a target performance metric ~~error rate~~ of a ~~every~~ zone (such as an intermediate zone) is not available~~,~~ then the histogram for that zone can be constructed by interpolation or ~~/~~extrapolation ~~is preformed to construct histograms for the intermediate zones~~. The histograms can be used to estimate a BPI distribution.

FIG. 2D shows a joint BPI distribution calculated from the histograms of the heads in the measured disk drives at a ~~The constructed histograms are used to calculate cumulative performance distribution functions (CDF) of the head frequency (e.g., BPI) capability at a given~~ target performance metric. The joint BPI distribution is a 2D distribution based on the histograms in FIG. 2C ~~error rate performance metric) as input to the format optimizer.~~ at the target on-track symbol error rate of $6x10^{-7}$. The x-axis is the BPI capability of the heads at the middle diameter (MD) of the disks, the y-axis is the BPI capability of the heads at the outer diameter (OD) of the disks, and the z-axis is the calculated number of heads divided by the total number of heads. The joint BPI distribution provides

an estimate of the probability that the heads meet the target performance metric at the MD and the OD.

The joint BPI distribution may predict, for example, that 10% of the heads in the measured disk drives can operate at a high frequency of 1.5 x the reference frequency, 50% of the heads can operate at a high frequency of 1.25 x the reference frequency, 90% of the heads can operate at the reference frequency, and 99.9% of the heads can operate at a low frequency of 0.75 x the reference frequency.

~~Such performance distribution functions are designated as marginal, individual or per zone distributions. In one example the distribution functions include one-dimensional (1-D), two-dimensional (2-D) and three-dimensional (3-D) joint BPI/frequency capability CDF, calculated at the same specified target error rate. Marginal or one-dimensional distribution functions from the joint distribution functions can also be calculated.~~

~~A version of estimating performance (e.g., frequency capability) cumulative distribution functions at a given desired target performance metric and zone is described. A number of frequency format profiles can be generated and tested on a set of disk drives to ensure proper operation. The frequency formats are generated and used to exploit every head's linear density or frequency/BPI sensitivity at every zone. Thus,~~ F~~f~~or example, the linear bit density sensitivity of every head at ~~a~~ z~~Z~~one K (where K ranges from 1 to M~~ zones~~) at the six frequency samples is determined. ~~To do so, the performance of every head is measured (after full drive read/write and servo calibration/optimization) at each frequency at ZoneK.~~ If frequency 1~~Freq1~~ K, frequency ~~Freq2~~ K . . . ~~, …,~~ f~~F~~requency 6 K are the frequency samples ~~selected frequencies~~ at z~~Z~~one K, ~~in the measurement process,~~ every head is positioned on ~~a track (e.g.,~~ the same track~~)~~ in ~~at~~ z~~Z~~one K and the record/playback performance of each head is

measured at every frequency sample using a target performance metric ~~of choice (e.g., off-track symbol error rate (SER))~~.

The BPI distributions can be calculated at the target performance metric as ID, 2D or 3D distributions that are marginal, individual or per zone distributions, respectively. The format optimizer uses the estimated frequency capability BPI distributions for every zone at the target performance metric to determine the storage capacity and yield.

~~For example, FIG. 2B shows the on-track SER performance of a selected head e.g. Head1, at Zone1, and the best least square polynomial fit curve. For a desired on-track target SER of 6e-7, the BPI capability of the (randomly) selected Head1 at that target performance metric can be projected by extrapolating the data (interpolation is performed if the desired on-track target SER is in the performance range, and extrapolation is performed otherwise). FIG. 2B shows the projected BPI capability of that head at an on-track SER of~~ ~~$6e^{-7}$ (i.e., $6x10^{-7}$).~~ ~~The original nominal kBPI (i.e., before the application of vertical zoning) is also shown. The head can be classified as a strong head because there is a reasonably significant amount of margin before the on-track SER performance of this head can be changed from its nominal frequency/kBPI of ~ 188 with an on-track (log of) SER of -9.1 to a projected on-track (log of) SER of -6.22 operating at a frequency/kBPI of ~ 217. Hence, there is a total kBPI gain of ~ 29, allowing increase in the nominal frequency by 15 % while meeting the desired target on-track SER performance metric of 6e-7 (i.e.,~~ ~~$6x10^{-7}$). Thus, for example, Head1 of disk drive3 has a frequency capability equal to 217 at an on-track SER of 6e-7 (i.e.,~~ ~~$6x10^{-7}$), which is noted for Head1 as one sample for generation and extraction of empirical histograms.~~

~~Performing the above steps for all the heads of all the disk drives considered in the measurement process, allows extraction of the empirical histograms of frequency capability at on-track SER of~~ ~~$6e^{-7}$ (i.e., $6x10^{-7}$) for such heads, as shown by example in FIG. 2D. The y-axis shows the number of heads~~

~~that meet the interpolated/extrapolated (i.e., projected) frequency capability that is shown on the x-axis. The extracted (empirical) histogram can be used to estimate the probability cumulative distribution function. The width of each histogram 76 corresponds to a variance of the head performance histogram, wherein an objective of the present invention is to improve the disk drive yield and capacity, and as a result reduce that variance.~~

~~FIG. 2C shows an example joint BPI distribution plot. Such BPI distributions may predict that, for example, 10% of the heads in the disk drives for which measurement was performed, can operate at a frequency density of 1.5 x freq (wherein freq is a reference frequency), and 50% can operate at density of 1.25 x freq, and 90% can operate at density of 1.0 x freq, and 99.9% can operate at density of 0.75 x freq, etc. Using the estimated frequency capability cumulative distribution functions at the target performance metric and every zone, the format optimizer determines the disk drive yield and capacity.~~

~~FIG. 2C is an example of a 2-D joint (i.e. outer diameter (OD) and middle diameter (MD)) BPI cumulative distribution function (CDF) at a target performance metric (e.g. on-track symbol error rate of 6e-7). The x-axis shows the BPI capability of all heads (i.e. from all the disk storage devices considered in the measurement phase) at MD AND y-axis is the BPI capability of all heads at OD. The z-axis shows the (calculated) number of heads divided by the total number/population of heads i.e., an estimate of probability that those heads have a joint BPI capability at OD AND MD of less than or equal to any given desired values. BPI capability for e.g. at OD in the above description it is meant that while operating at (or below, i.e., if considering CDF) a given BPI, after full Read/Write and servo calibration and optimization, meeting a desired given target performance metric of choice at OD. Thus, the aforementioned description of BPI capability can similarly be extended to joint BPI capability.~~

Format Optimizer ~~process~~

The format optimizer ~~process block~~ 66 provide~~comprise~~s ~~a~~ variable BPI optimization. The format optimizer 66 solves ~~process for solving multiple (e.g., three)~~ constrained optimization problems in response to various inputs. ~~given the inputs: number of frequency formats (i.e., desired number of different read/write frequencies), number of heads in each disk drive, and~~ the~~said BPI distributions.~~ The first problem maximizes the ~~drive~~ yield while preserving the storage ~~same drive~~ capacity, the second problem maximizes the storage ~~drive~~ capacity while preserving the ~~same drive~~ yield, and the third problem maximizes the ~~drive~~ yield while reducing the track density and meeting the storage capacity~~allowing reduced/relaxed TPI such that the desired target drive capacity can be met~~. The inputs include the number of different read/write frequencies (frequency profiles or formats), the number of heads in each disk drive, the BPI distributions, and the nominal storage capacity. The BPI distributions indicate the frequency capability distribution of the heads at a target performance metric.

T

~~The format optimizer 66 inputs~~ the~~said BPI distributions, a desired disk drive capacity, the total number H of heads per disk drive and the number N of frequency profiles or vertical zones (i.e., frequency per zone across a~~ disk~~media surface stroke). Given the frequency capability distribution of heads at a target performance metric, the total number of vertical frequency formats (F), the total number of heads per drive (H) and the nominal drive capacity,~~ the format optimizer 66 simultaneously searches through a~~all possible~~ continuous range of all possible frequency capabilities to maximize the ~~drive~~ yield such that the ~~desired~~ nominal storage ~~drive~~ capacity is met. The format optimizer 66 can also perform~~solve~~ the same operation with the ~~problem, but the drive~~ storage capacity and the ~~drive~~ yield interchanged ~~interchanged~~.

The ~~As such, in one example, the~~ format optimizer 66 can optimize~~s~~ high and ~~/~~low data density as a function of the zones~~, to improve the drive yield and meet a fixed target capacity. The format optimizer also optimizes capacity while achieving a fixed nominal drive yield, wherein the nominal yield is before the application of vertical zoning according to the present invention~~. For~~In~~ the~~said~~ example, in a disk drive with eight~~of F=2 formats (high density and low density) and H=8~~ heads ~~16 in a disk drive 100~~, the possibilities are one~~include 1~~ head at high data density and seven~~7~~ heads at low data density, two~~2~~ heads at high data density and six~~6~~ heads at low data density, three~~3~~ heads at high data density and five~~5~~ heads at low data density, four4 heads at high data density and four~~4~~ heads at low data density, one~~1~~ head at low data density and seven~~7~~ heads at high data density, two~~2~~ heads at low data density and six~~6~~ heads at high data density, and three~~3~~ heads at low data density and five~~5~~ heads at high data density~~, etc~~. T~~Hence, t~~he format optimizer 66 considers all the combinatorial possibilities, ~~and~~ in each case solves a constrained optimization problem and ~~finally~~ chooses the ~~best~~ optimal solution among~~st all~~ the possibilities. Alternatively, the format optimizer 66 can ~~be designed to~~ reach the ~~best~~ optimal solution more directly by ~~solving a~~ ~~(~~non-linear~~)~~ mixed- integer programming.

Therefore, once the 1-D, 2-D and 3-D ~~joint~~ BPI distribution~~/frequency capability CDF~~s ~~(discussed above)~~ at a ~~given~~ target performance metric are ~~calculated and~~ passed ~~as input~~ to the format optimizer 66, the format optimizer 66 solves ~~the above~~ two problems, ~~namely:~~ (1) maximizing or improving the ~~drive~~ yield ~~(i.e.,~~ due to the target ~~pre-selected~~ performance metric~~, e.g., off-track SER)~~ while meeting the ~~a~~ desired nominal storage ~~drive~~ capacity, and (2) maximizing the storage ~~disk drive~~ capacity while meeting the~~a~~ desired nominal ~~drive~~ yield.

The format optimizer 66 ~~then~~ mathematically casts these two problems ~~stated above~~ as constrained optimization problems and solves them using well-known optimization techniques such as a ~~e.g.~~ line search algorithm. The constrained optimization problems can also be cast as ~~(~~non-linear~~)~~ mixed-integer

programming and solved using existing ~~methods in~~ optimization methods~~theory~~. Example constraints to be considered, and cast mathematically within the format optimizer 66, include not exceeding a certain frequency at the outer diameter ~~OD~~ due to ASIC data rate limitations or at the inner diameter ~~ID~~ due to head/disk~~media~~ limitations. Furthermore, closed form equations are derived and used in the format optimizer 66 to estimate the storage capacity and ~~actual drive~~ yield~~and drive capacity~~. The format generator 66 also considers ~~A Format Generator process, described below, is utilized to calculate the actual drive capacity after including all~~ possible overheads~~,~~ such as adding ~~and including~~ redundant bits due to error correction coding or gray coding.

The format optimizer 66 also uses ~~the~~ information from the format generator 68~~,~~ such as the calculated format efficiency per zone (~~i.e.,~~ defined in percentages as the amount of user data ~~e.g.~~ in blocks that can fit in all tracks in a zone), or the number of tracks per zone, to achieve a very close estimate of the storage ~~disk drive~~ capacity ~~calculation as~~ determined by the format generator 68. Then, the format optimizer 66 calculates optimal linear bit density format profiles as well as the optimal number of heads allocated to each vertically zoned format profile.

For ~~As an~~ example, histograms are extracted and the corresponding BPI distributions are estimated for different zones at the ~~desired~~ target on-track symbol error rates ~~(e.g., for Zone1 at a target error rate of $1e^{-6}$, Zone2 at a target error rate of $1e^{-6}$, Zone3 at a target error rate of $1e^{-6}$, etc.) as described above~~of $6x10^{-7}$. A format design is provided for a disk drive with four4 -heads ~~disk drives~~ ~~(H~~and two frequencies to optimize yield ~~=4), and 2 vertical frequency format/profiles (F=2), wherein the disk drive yield is optimized~~ while meeting ~~a minimum~~ storage capacity~~requirement~~.

~~Without the method of the present invention (vertical zoning), conventionally when the same frequency is used per head per zone, if one of the~~

~~4 heads is a weak performing head having an error rate measurement e.g. 1e^{-5} at Zone1 (higher than the target error rate), that disk drive is failed. With the application of the present invention in that case, the format optimizer allocates the 3 other heads to higher frequencies, and allocates the weak head to a lower frequency at Zone1. The recording/playback performance of the weak head is compensated for, such that the minimum capacity requirement is met. As such, the format optimizer 66 utilizes the performance distributions to determine two or more optimal frequencies per zone, and the optimal number of head allocations to those frequencies per zone such that constraints such as required disk drive yield and/or capacity are met.~~

For example, the format optimizer 66 uses the ~~estimated~~ 1D, 2D and 3D ~~joint frequency/~~BPI ~~capability~~ distributions at the ~~a desired~~ target performance metric to jointly optimize for vertically zoned frequency format profiles and the corresponding number of head allocations three zones at a time. An advantage of considering three zones instead of one zone,~~, as compared to only one~~ ~~(~~and thus ~~considering a~~ joint optimization instead of~~versus~~ individuali~~zed~~ optimization,~~),~~ is that the joint optimization allows the ~~optimization of format profiles (e.g.,~~ frequency profiles to be optimized~~)~~ across the stroke on each disk surface. _Therefore, joint optimization ~~in this way we~~ exploits the potential correlation in performance from one zone to another zone as well as their individual and weighted contribution to the storage ~~overall surface~~ capacity. JA joint optimization ~~approach~~ is preferable for a high/low/~~high~~ data density format ~~layout~~ across the stroke for either improving the ~~drive~~ yield while keeping the same storage~~drive~~ capacity~~,~~ or improving the storage ~~drive~~ capacity while preserving the ~~same drive~~ yield.

The format optimizer 66 generates~~d results from the format optimizer include:~~ the target high/low BPI formats per zone, the ~~(~~optimal~~)~~ number of head allocations per format,~~/layout~~ and an estimate of the storage capacity and ~~drive~~ yield~~ and capacity~~. -The accuracy of the estimates can be sensitive to the

underlying ~~extracted (joint)~~ BPI distributions at the target performance metric~~a given (on or off-track) error rate~~. Further, the target high/low BPI formats can be sensitive to the variance of the ~~(extracted)~~ BPI distributions. And, the variance of the BPI distributions can be sensitive to the absolute value of the target performance metric and the type of target performance metric~~(on or off-track) target error rate and the choice of metric (e.g. error rate vs. MSE)~~. In addition, ~~because~~ the ~~design of~~ target high/low BPI formats are designed ~~performed~~ three zones at a time and the yield improvement ~~(~~while preserving the storage ~~same overall drive~~ capacity~~)~~ is based on the profile of the target nominal formats. ~~,~~ T~~t~~he format optimizer 66 also allows for ~~a~~ smoothing ~~operation in order to smooth~~ the target variable BPI format designs~~if so desired~~. The format generator 68 determines the number of tracks per zone, the number of blocks per track, the radius at each zone, as well as block and track format efficiency. This information is saved in ~~e.g.~~ output files for use with the format optimizer 66. The format optimizer 66 then saves the ~~design of~~ target high/low BPI formats per zone that it generates~~,~~ in two separate files that can be ~~read and~~ loaded ~~as input files~~ into the format generator 68.

Once the target format profiles are calculated, if they are non-smooth across the stroke, optionally a smoothing process is applied. The format profiles are then loaded into the format generator 68 ~~described below,~~ to create vertically zoned formats and configuration pages. The formats and configuration pages are used by the disk drive firmware to create binary files to be loaded in~~on~~to the reserve image of the disk drives as part of the file system. ~~FIG. 2A shows the communication between the format optimizer 66 and the format generator 68.~~ In this fashion, the design and implementation of the format profiles~~,~~ as well as the number of optimal head allocations are performed off-line and are predetermined for every disk drive configuration.

For ~~An~~ example, in a disk drive with four heads and four disk surfaces on two disks, the format optimizer 66 designs ~~of format optimization for designing~~

vertically zoned high and low ~~and high~~ frequency profiles~~ for disk drives with 4 heads (i.e., Head1 through Head4 corresponding to Surface1 through Surface 4 of Disk1 and Disk2) is described~~. E~~In this example, e~~very disk surface is ~~(~~uniformly~~)~~ partitioned into three zones across the stroke, at a track density with a fixed number of tracks ~~(TPI)~~ per zone, vertically aligned from one disk surface to another. The nominal disk surface data storage~~capacity~~ ~~(~~before the ~~application of~~ vertical zoning~~)~~ can be approximated by the sum over all the zones of the ~~products of~~ nominal tracks per zone multiplied by the nominal BPI ~~frequency~~ per track~~zone~~ multiplied ~~and~~ by the format efficiency per zone. ~~-~~Format efficiency per zone is the~~a~~ percentage of ~~all~~ the user data that is effectively stored per zone. T~~Then~~ the nominal storage ~~disk drive~~ capacity is ~~equal to~~ the nominal disk surface data storage~~capacity~~ multiplied by the total number of disk surfaces (or heads). The nominal number of tracks per zone and the format efficiency per zone can be generated by ~~and obtained from~~ the format generator 68~~(described further below)~~.

Performing vertical zoning to ~~e.g.~~ improve the ~~drive~~ yield without losing a storage ~~ny nominal (disk drive)~~ capacity~~,~~ finds the best frequency per zone and per head such that the disk drive meets performance and storage capacity requirements. I~~As such, i~~f a disk drive with four~~4~~ heads fails ~~test process performance limits~~ due to the ~~e.g.~~ performance of h~~H~~ead_1 at z~~Z~~one_1, but the performance of another head/zone pair, such as h~~H~~ead_1 at z~~Z~~one_2 ~~(~~or ~~another head such as~~ h~~H~~ead_3 at z~~Z~~one_1,~~)~~ ~~performance~~ is significantly better, ~~(i.e.,~~ passing the test~~ limits~~ with reasonable margins~~)~~, then a higher than nominal frequency ~~than nominal~~ at z~~Z~~one_1~~2~~ or z~~Z~~one_2~~1~~ is designed for the strong heads~~ that are stronger (i.e. high density heads),~~ and ~~instead~~ the frequency at z~~Z~~one_1 for the weak heads ~~that are weaker (i.e. low density heads)~~ is lowered. This trade-off ~~is performed~~ obtains a vertically zoned design of variable frequencies per zone such that the storage ~~overall disk drive~~ capacity is preserved~~, to obtain a vertically zoned design of variable frequencies per zone~~. In addition, the number of heads ~~(e.g.~~ per zone~~)~~ allocated to high or low data

density is determined. _Thus, the storage ~~disk drive~~ capacity can be approximated by the sum ~~(~~over all the zones~~)~~ of ~~the products of~~ the number of strong ~~low density~~ heads ~~per zone~~ multiplied by the high~~low~~ frequency data storage per zone multiplied by the format efficiency per zone plus the ~~(~~sum over all the zones of ~~the) products of~~ the number of weak ~~high density~~ heads ~~per zone~~ multiplied by the low~~number of high~~ frequency data storage~~heads~~ per zone multiplied by the format efficiency per zone.

For example, t~~In one version, the format optimizer solves for the above problem as follows. T~~he format optimizer 66 is provided with joint ~~the~~ BPI ~~(joint) frequency capability cumulative~~ distribution~~-functions~~ ~~(extracted and estimated from all heads considered in the measurement process above)~~ at the ~~a desired~~ target performance metric~~(i.e., the same targets used in the test process)~~. Then, for every combinatorial possibility of head allocation ~~(e.g.~~ to high or low frequency,~~)~~ the format optimizer 66 searches through a continuous range of possible frequencies~~,~~ by considering every zone independently ~~(i.e.~~ using the marginal distributions~~)~~ and by the combination of zones~~(i.e.~~ using the joint BPI distributions~~),~~ to maximize the ~~disk drive~~ yield calculated using a closed form equation, such that the storage ~~disk drive~~ capacity after the ~~application of~~ vertical zoning is applied is essentially the same as the nominal storage ~~disk drive~~ capacity. Further, the ~~(~~optimal~~)~~ high and low ~~and high~~ frequency profiles for every combination of head allocations is compared and the one that results in the highest ~~value of (calculated) disk drive~~ yield is chosen and passed to the format generator 68 for the generation of vertically zoned configuration pages to be used by the disk drive firmware.

The disk surfaces ~~of disks~~ can be partitioned into more than ~~the example~~ three zones~~above~~. ~~The above steps of determining the optimal variable frequencies per zones are useful to consider more than three zones.~~ To reduce computational complexity and time, if the selected/designed number of zones per disk surface ~~in a disk drive~~ is more than three, the format optimizer 66 can ~~be~~

~~used to~~ generate high and low ~~and high~~ frequency profiles three zones at a time and~~,~~ ~~suitable smoothing operations are used to~~ smooth the profile after post-processing. Another approach includes embedding the smoothing operator in the design and extending the joint optimization to all the zones ~~so as~~ to consider the ~~effect and~~ impact of smoothing to ~~drive~~ yield ~~(~~calculation~~)~~ as part of the design rather than the later stages.

In the disk drive with four ~~above 4~~-heads~~-disk drive example~~, the ~~disk drive~~ yield is maximized while preserving the ~~same~~ nominal storage ~~disk drive~~ capacity. To determine the number of head allocations, the format optimizer 66 begins with one weak ~~one low density~~ head and three strong ~~high density~~ heads ~~(e.g.,~~ per zone~~)~~. The format optimizer 66 searches through a continuous range of possible frequency capabilities per zone, as well as two and three zones at a time, by considering ~~and using~~ the 1D ~~(i.e., marginal)~~, ~~(joint)~~ 2D and 3D BPI distributions that result in the best calculated ~~value of the drive~~ yield such that a minimum nominal storage ~~drive~~ capacity can be obtained. Next, the format optimizer 66 uses two weak heads and two strong ~~low and high density~~ heads and repeats solving the ~~same~~ constrained optimization problem. This process is continued until all the combinatorial possibilities are considered. Finally, the format optimizer 66 chooses the solution that results in the best ~~calculated value of the drive~~ yield and provides the target high and low ~~and high (~~optimal~~)~~ data density format profiles and the associated number of high and low head allocations to the format generator 68. The format generator 68 then generates vertically zoned format files and configuration pages to be used by the disk drive firmware.

Format Generator

~~In one embodiment, the format generator process block 68 is used for embedded servoing (i.e., servo position is generated by reading back written information from the disks, such as servo wedges in which position information is~~

~~embedded on the disks, and that information is used to position the head on the disk surfaces).~~

T~~For example, t~~he format generator 68 ~~described herein~~ generally performs three functions. First, the format generator 68 uses~~ including utilizing~~ target formats/frequencies (or linear densities/BPI) for each zone ~~as an input,~~ and ~~performing an exact~~ calculates~~ion~~ ~~of~~ the data storage ~~capacity~~ of each zone and thus the storage capacity of the disk drive~~ itself~~. Second~~Further~~, the format generator 68 calculates the format efficiency (the percent of the disk surface ~~area~~ that is occupied by user ~~customer~~ data) for each zone. Third,~~The third, and primary purpose of~~ the format generator 68 ~~is to~~ generates configuration pages ~~data~~. The configuration pages contain per-drive, per-zone, and per-head-per-zone parameters that are programmed into the disk drive electronics such as ~~components. Such components include~~ the preamplifier 21, the ~~disk controller, the~~ read/write channel 51~~,~~ and the controller 57~~preamplifier~~. The parameters are ordered such that the disk drive firmware selects the correct set of parameters to be programmed into each of the components for the particular head and zone that is being written to or read from at the time.

The format generator 68 calculates the ~~exact~~ frequency and the data storage ~~exact capacity~~ of each zone taking into consideration limitations in the programmability and ~~of the components and limitations of~~ the capability~~ies~~ of the disk drive components. For example,~~Some examples of component limitations include:~~ the heads 16 have varying down-track separation between the read and write elements, the preamplifier 21 has a minimum and maximum delay in turning on the write current, the read/write channel 51 synthesizer frequencies are limited to ~~a~~ discrete ~~collection of~~ frequencies, the motor driver 53 can keep the spindle motor 14 within a finite precision of the nominal rotational speed,~~;~~ ~~preamplifier (not shown) has a minimum and a maximum delay in turning on its write current; the down track separation between the head read and write elements (not shown) varies from component to component; the reference~~

~~crystal (not shown) has finite accuracy and stability over temperature; the spindle motor driver can keep the spindle motor speed within a finite precision of the nominal rotational speed;~~ the controller 57 has specific latencies in generating commands to the preamplifier 21 and the read/write channel 51~~and preamplifier,~~ often with a finite uncertainty as to the exact timing of these commands, and a reference crystal (not shown) has finite accuracy and stability over temperature~~, etc~~.

The format generator 68 can be fully automated, or can be directed by a human operator~~specialist~~. In the absence of input from the format optimizer 66, the target per-zone BPI/frequency profiles, in particular, must be generated by~~from~~ a human operator~~input~~. In general, the human operator~~specialist~~ modifies the target frequency profiles until the desired storage capacity is reached.

T~~In one embodiment, t~~he format generator 68 includes a format efficiency process that uses the format optimizer 66~~'s~~ target high/low~~/high~~ variable BPI format designs as well as the optimal predetermined number of high/low~~/high~~ performing head allocations~~,~~ to modify and generate the appropriate configuration pages ~~(i.e.,~~ as part of the file system~~)~~. For each zone, the format generator 68 selects the nearest frequency to the target frequency for that zone, given the component limitations ~~and programmability~~ mentioned above. The nearest frequency provid~~comprise~~es the target formats.

The optimal predetermined number of ~~low/~~high/low performing head allocations comprises the number of heads allocated to each of the multiple frequencies in each zone. The format optimizer 66 determines the head allocation, which is input to the format generator 68. The capacity of a zone depends ~~both~~ on the target frequencies and the number of heads allocated to each frequency.

The format optimizer 66 uses the nominal average BPI or frequency (nominal BPI format target designs) (e.g., one read/write frequency) in each zone ~~as input~~ from the format generator 68 to estimate the ~~disk drive~~ yield before applying the variable BPI designs. For a design with multiple frequencies per zone, this is the weighted average ~~(~~by the number of allocated heads~~)~~ of the multiple frequencies. The nominal format is created by ~~e.g.~~ a human operator working with the format generator 68 in an~~the~~ interactive manner~~described above~~.

The format generator 68 ~~performs~~ calculate~~ions~~ the~~of~~ number of tracks per zone, number of blocks per track, radius at each zone as well as block and track format efficiency to calculate the ~~drive~~ zone data storage~~capacity~~. -The format optimizer 66 estimates the zone data storage ~~capacity~~ using the tracks per zone, radii, and format efficiency. Thus, the format optimizer 66 and the format generator 68 interact as shown in FIG. 2A~~6~~. For example, in a disk drive with four ~~for 4~~ heads~~ disk drives~~, and two~~2~~ data density format frequency profiles (~~i.e.,~~ high and low frequency profiles) with three~~3~~ zones across the disk surface, after the measurement and optimization processes, the format generator 68 is provided with two ~~2~~ optimal frequency profiles and the optimal allocation of the heads. The format generator 68 then calculates the storage capacity, and if the disk drive ~~capacity~~ meets the minimum required storage~~required~~ capacity, ~~then~~ the format generator 68 generates the configuration ~~files/~~pages for the disk drive firmware. The configuration pages are used by the disk drive firmware to command the head to write at an assigned frequency to a zone. If the calculated storage ~~drive~~ capacity does not meet the minimum required storage capacity, the format optimization is performed again with new format efficiency values~~,~~ and the process is repeated.

Head ~~format~~ Aassignments~~ and selection criterion~~

Allocating ~~The process for allo~~the ~~cation of~~ number~~s~~ of heads to ~~each of~~ the predetermined multiple frequencies in a zone, and assigning ~~the process of~~

~~assignment of~~ a particular head in a particular disk drive to a particular frequency, are distinct. T~~First t~~he allocation ~~process~~ is- performed by the format optimizer 66~~,~~ and applies to the disk drives of a particular design~~(product)~~. Then~~, the~~ head assignments are~~process is~~ then performed during manufacturing as part of a test process undergone by each disk drive to be produced.~~ This section describes the assignment process task.~~

Once the configuration pages are generated and converted to binary files as part of the file system, they can be loaded ~~and saved~~ into a reserved image of the disk drive for use after power cycling. Then, for every disk drive, the ~~following example~~ assignments are~~process is~~ performed per head and per zone~~,~~ to ~~determine~~ assign~~ment of~~ a ~~certain~~ predetermined number of heads to high BPI formats and the remaining heads to low BPI formats in a two~~2~~ frequency design, to satisfy the allocation of heads to the~~said~~ formats by the format optimizer 66.

The head assignments ~~process~~ for the two ~~example 2~~ frequency format where high and low frequencies are used~~,~~ includes the steps of:

(1)~~.~~ Load default parameters from the configuration pages, and calibrate selected parameters on a per head, per zone basis (e.g., load high BPI format profile for all the zones across the stroke);

(2)~~.~~ Take measurements from the~~all~~ heads at ~~all~~ the disk surfaces at pre-selected zones with respect to a target performance metric~~, e.g., mean square error, on/off track error rate, etc.~~;

(3)~~.~~ For each head in every measured zone, sort/rank the heads by the target performance metric from best to worst,~~;~~ select a pre-specified (by the allocation process in the format optimizer 66) number of heads with the best performance, and assign those heads to the high~~er~~ frequency for a particular zone;

(4)~~.~~ Optionally interpolate between the measurements obtained from the pre-selected number of zones to find the results for the other zones, and do the same for the interpolated zones. The interpolation ~~operation in a version of~~

~~the present invention~~ reduces the test ~~process~~ time. Head performances are measured, sorted and assigned to a frequency for a subset of the total number of zones. For the remaining zones, the heads are assigned by interpolating ~~on~~ the head assignments ~~made~~ from the measurements;

(5)~~.~~ For every zone, save the worst pre-specified number of weak ~~bad (i.e. low performing)~~ heads with respect to the target performance ~~selected~~ metric; and

(6)~~.~~ For every zone, load and calibrate ~~all~~ the weak ~~bad~~ heads with the low~~er~~ BPI format.

The above process can ~~be used to~~ improve storage ~~yield, improve~~ capacity, improve yield and trade-off between ~~yield and~~ storage capacity and yield. In a test, the heads can ~~be~~ pass~~ed~~ or fail~~ed~~ with respect to a target performance metric ~~(e.g., off track error rate)~~ to determine if the test target limits are met.

The disk drive~~servo~~ firmware is extended to load more than one format profile. A head can be assigned a different read/write frequency per zone across a disk surface, and radially similarly situated zones on different disk surfaces can have different read/write frequencies assigned to the corresponding heads whereby one head is assigned a different frequency/format profile than another head.

The head ~~example a~~assignments ~~process described herein~~ apply~~ies~~ to a format design with two recording frequencies per zone, but~~. However, the process~~ can be easily extended to more than two frequencies per zone and ~~,~~ ~~wherein the process~~ can be iterated ~~upon~~ to assign heads to more than two frequencies per zone~~, as described by an example below~~. For example, in a design with H heads and F frequencies per zone, ~~above~~ steps 1 and 2 are completed for the high frequency. The first selection of heads in step 3 assigns the highest h1 heads, where h1 is the pre-specified number of heads allocated to

the highest frequency for that zone. The remaining (H -- h1) heads are then loaded and calibrated with the second highest frequency (step 1 again), measurements are taken (step 2 again), ~~and~~ the heads are ordered relative to the metric and the best h2 heads are assigned to the second highest frequency (step 3 again). Here h2 is the pre-specified number of heads allocated to the second highest frequency in the zone. Steps 1-3 ~~of the process~~ are then iterated for the (H -- h1 -- h2) heads, followed by the (H -- h1 -- h2 -- h3) heads, and so on, until hF heads remain to be assigned to the lowest frequency. -The set of {h1 . . . ~~, ...,~~ hF} heads receive ~~comprise~~ the head allocation made by the format optimizer 66.

Table 1 ~~below~~ illustrates ~~the result of an example of~~ the vertical zoning head assignments ~~process~~ on a disk drive with six~~6~~ heads and five~~5~~ zones across the stroke on each disk surface. Each head is assigned to either a high or low data density format based on record/playback performance of that head, and ~~wherein as discussed above,~~ the number of heads assigned to high data ~~low~~ density and the number of heads assigned to low data ~~high~~ density is according to the head allocation ~~results~~ determined by the format optimizer 66.

HEAD #	~~FORMAT\ZONE~~ ZONE 1	ZONE 2	ZONE 3	ZONE 4	ZONE 5
0	Low	High	Low	High	Low
1	High	Low	High	High	Low
2	High	Low	High	Low	High
3	High	High	Low	High	High
4	Low	High	High	High	High
5	High	High	High	Low	High

Table 1 – E~~An~~ example ~~for the~~ format assignment of a disk drive after test ~~process,~~ using vertical zoning with variable BPI across ~~the~~ zones.

~~FIGs. 3-6 show example steps of an embodiment of the above processes. Referring to~~ FIG. 3 shows~~,~~ a flowchart of ~~n example~~ vertical zoning data collection that~~process~~ includes the steps of:

(1) Select a number of disk drives for data measurement/collection ~~process~~ (step 300);

(2) Create a nominal linear bit density profile KFCI (~~i.e.,~~ nominal_KFCI):

$\overline{\text{kFCI}}(R)$, where~~in~~ R is the disk radius (step 302);

(3) Create more linear bit density profiles by multiplying the nominal_KFCI by the scaling factor x_i~~,~~ (step 304):

$$(1 \pm x_i) * \overline{\text{kFCI}}(R)$$

(4) Create a binary file system for every generated profile (step 306)~~,~~ i~~.e. for~~:

$$i \in \{1, \cdots, N\}$$

where~~in~~ N is the total number of frequency format profiles. ~~-~~ f~~F~~or example, for N = 2, having X_1, and X_2, if X_1 = 0.05 and X_2 = 0.1, then including the nominal frequency format ~~itself,~~ there are five~~5~~ different frequency profiles in step 304~~,~~ as follows: (a) nominal_KFCI, (b) 1.05 x ~~*~~ nominal_KFCI, (c) 0.95 x ~~*~~ nominal_KFCI, (d) 1.1 x ~~*~~ nominal_KFCI, and (e) 0.90 x ~~*~~ nominal_KFCI ~~(wherein "*" is the multiplication operator)~~;

(5) Select the first head by setting *i* to 1(step 308);

(6) Load the file system *i* i~~o~~nto the reserved image of the disk drives (step 310);

(7) Take the head performance measurements ~~(e.g., (on/off) track MSE and SER measurements)~~ (step 312);

(8) Unload and save the results in the d~~D~~ata b~~B~~ase (step 314);

(9) Increment i by one~~ (i=i+1)~~ (step 316);

(10) Determine if ~~is~~ i = N~~?~~ -(step 318);

(11) ——————If not, go to step -310; and

(12) Otherwise, ~~else~~ stop (step 320)~~done~~.

The above process collects performance data for all the heads at all the zones.

~~Referring to~~ FIG. 4 shows~~,~~ a flowchart of ~~n example~~ vertical zoning post-measurement ~~processing~~ and per zone BPI distribution extraction that ~~process~~ includes the steps of:

(1) Organize the head performance data ~~(e.g., MSE and SER) obtained above,~~ for every head $i \in \{1, \cdots, M_1\}$ and every zone $j \in \{1, \cdots, M_2\}$ as a function of the linear bit density samples, where~~in in this example~~ M_1 is the total number of heads in the disk drives selected for ~~the~~ measurement ~~process (e.g., 40 disk drives each including 4 heads, results in total of M1=160 heads),~~ and M_2 is the total number of zones, to generate head performance histograms (step 400);

(2) Choose a target performance metric ~~(e.g., MSE or SER)~~ (step 402);

(3) Set $j = 1$ and $i = 1$ (step 404);

(4) Interpolate/e~~E~~xtrapolate BPI at the ~~specified~~ target performance metric ~~(e.g., MSE or SER)~~ for head i at zone j -(step 406);

(5) Select the next head by incrementing i by one ~~(i = i +1)~~ (step 408);

(6) Determine if ~~Is $i = M_1$? (i.e., have~~ all the heads have been processed~~?)~~ by determining if $i = M_1$ (s~~S~~tep 410);

(7) If not, go~~t~~ to step 406 to process the next head, otherwise~~else~~ generate a frequency capability histogram at ~~that given~~ zone *j* for all the heads (step 412);

(8) Determine if ~~Is $j = M_2$? (i.e., have~~ all the zones have been processed~~?)~~ by determining if $j = M_2$ (s~~S~~tep 414);

~~(9) If yes, stop;~~

(9~~10~~) If not, ~~Otherwise,~~ move to the next zone and start with the first head again, set ~~whereby~~ $j = j + 1$ and~~,~~ $i = 1$ (step 416)~~,~~ and go to step 406; and

(10) Otherwise, stop (step 418)~~to repeat~~.

The above process generates ~~(~~1D~~)~~ frequency capability histograms at a target performance metric ~~and~~ for every zone by considering all the heads from the sample disk drives selected for ~~in the~~ measurement~~ process~~. Using the ~~(~~1D~~)~~ frequency capability histograms at a ~~given~~ target performance metric, ~~techniques in~~ probability theory known to those skilled in the art can be adopted to estimate the ~~(~~1D~~)~~ frequency capability distributions. Further, the above process~~dure above~~ is extended ~~(i.e.,~~ by using 2D and 3D interpolation/extrapolation routines~~),~~ to extract and estimate the 2D and 3D joint frequency capability histograms and their associated distributions.

~~Referring to~~ FIG. 5 shows~~,~~ a flowchart of head assignments for N heads with ~~n example head assignment process for~~ a two frequency format (~~N=2,~~ high/low data density) that ~~design,~~ includes the steps of:

(1) Assign all~~n~~ the heads in a disk drive to the ~~first selected format (e.g.,~~ high data-density format~~)~~ (step 500);

(2) Calibrate all~~n~~ the heads at the high data-density format for selected zones (step 502);

(3) Measure the head performance metric at the selected zones for all the~~n~~ heads (step 504);

(4) For each selected zone, rank the heads by the head performance metric (step 506);

(5) For each selected zone, assign the highest K heads to the high data ~~-~~density format, and assign the other N~~n~~-K heads to the low data density format (step 508);

(6) Optionally interpolate the head assignments for the remaining zones (step 510); and

(7) Complete the calibration of all the heads and all the zones at the assigned formats (step 512).

The above process completes the assignment of each head in each disk drive to a predetermined frequency.

T~~As shown in FIG. 2A, information is passed between the format generator 68 and the format optimizer 66, wherein initially, t~~he format generator 68 passes the ~~information including e.g.~~ track formats~~layout~~ to the format optimizer 66 to have a more accurate way of calculating the storage capacity (nominal format). Such information and constraints are provided to the format optimizer 66 to solve the~~said~~ joint optimization problems. The format optimizer 66 performs a coarse -calculation of the storage capacity, whereas the format generator 68 performs an exact calculation of the storage capacity. The format generator 68 ~~performs functions of~~ provides~~ing~~ format information (such as ~~e.g.,~~ number of tracks per zone, and the zone format~~layout~~) to the format optimizer 66, and calculates~~ing~~ the exact storage~~format~~ capacity. Such information is passed once from the format generator 68 to the format optimizer 66 for a head design ~~(e.g., 4 head design)~~ with a given number of heads. -The format generator 68 initially provides nominal information- to the format optimizer 66, and ~~wherein~~ the format optimizer 66 performs its calculation of target densities (zone frequencies and number of heads allocated to each frequency) and provides that information

to the f~~F~~ormat ~~-~~g~~G~~enerator 68. The format generator 68 then determines if required storage capacity has been reached. Adjusting the target densities to meet storage capacity and/or yield ~~and/or capacity~~ requirements includes adjusting the selected zone density or zone frequencies.

~~Referring to~~ FIG. 6 shows ~~,~~ a flowchart of format generation and optimization in which an i~~n example Format Generator/Format Optimizer~~ ~~I~~terative~~eon~~ process for a minimum storage capacity ~~requirement~~ (C)~~,~~ and a user specified storage ~~allowed~~ overcapacity ~~Delta,~~ (Δ) includes the steps of:

(1) Determine the disk geometry, track density ~~(TPI)~~ and servo spoke~~wedge~~ details, and provide the ~~output~~ inner diameter (ID) and outer diameter (OD) radii, ~~;~~ the track density (TPI), ~~;~~ the number of servo spok~~wedg~~es, and the servo spoke~~wedge~~ length (step 600);

(2) The f~~f~~ormat generator 68 generates the initial format at the storage capacity using the ~~values~~ ID and ~~,~~ OD radii, the TPI, the number of servo spok~~wedg~~es and the servo spoke ~~wedge~~ length, and provide~~outputs~~ the radius of each zone per disk surface, the number of tracks per zone, the number of blocks per track, and the format efficiency by zone (step 602);

(3) The f~~f~~ormat optimizer 66 generates optimal target densities at all the zones using the radius of each zone per disk surface, the number of tracks per zone, the number of blocks per track, and the format efficiency by zone~~as described above,~~ and provides the high and low BPI ~~outputs frequency density~~ targets ~~(e.g., low/high BPI)~~ by zone~~,~~ and the number of high and low BPI ~~frequency density (e.g., low/high BPI)~~ head allocations by zone (step 604);

(4) The f~~f~~ormat generator 68 generates new formats with a storage c~~C~~apacity (the~~i.e.,~~ number of logical blocks per disk drive) (step 606);

(5) Determine i~~I~~f the storage c~~C~~apacity > C and the storage c~~C~~apacity < (C + Δ~~Delta~~) (step 608);

(6) If not, ~~, then stop;~~

~~(6) Otherwise,~~ adjust the target densities (s~~S~~tep 610)~~,~~ and go to step 606; and

(7) Otherwise, stop (step 612).

For ~~In one~~ example, the disk surface capacity is ~~described by the equation:~~ TPI x BPI x (1 + ECC) / FE, where~~in~~ TPI is the track density, BPI is the linear bit density, ECC is the fractional level of error correcting code ~~used~~ which is typically about 0.1, and FE is the format efficiency which is typically about 0.57.

The above process completes the format generation ~~process~~.

As another example ~~scenario of the results generated by an embodiment of vertical zoning according to the present invention~~, a set of thirty-two~~32~~ mature~~d~~ disk drives are selected and ~~wherein~~ each disk drive includes twelve~~12~~ heads. The 1D, 2D and 3D ~~joint~~ BPI ~~empirical~~ distributions are extracted at an ~~given specified target~~ on~~-~~track symbol error rate from the ~~three pre-specified radial zones, i.e.,~~ outer, middle and inner zones. Next, the BPI ~~extracted~~ distributions are fed ~~in~~to the format optimizer 66, and high or low ~~or high~~ frequency per zone format designs are obtained at the three ~~specified~~ zones. This is performed once by individual optimization~~, all~~ based on 1D BPI distributions at each of the three zones, and once by joint optimization based on the measurements obtained from the three zones and their extracted 1D, 2D and 3D BPI distributions. The head format allocation search ~~process (VZ test)~~ is performed by~~in a~~ simulation, and ~~wherein~~ for each zone the one-format designs ~~(i.e.~~ before the application of vertical zoning~~)~~ are a special case of the two-format variable BPI designs by forcing the high and low ~~and high~~ formats to be ~~the same and~~ equal to the nominal BPI format at that zone. Furthermore, the pass/fail of the disk drives is ~~decided~~ based on ~~the criterion that~~ each head at every zone passing a target ~~given~~ on-track symbol ~~target~~ error rate as well as off-track squeeze and un-squeeze offset margins. Then, the ~~drive~~ yield is calculated by~~(i.e., in~~ simulation by interpolation/extrapolation of the measurement data~~)~~ before and after the application of vertical zoning (VZ). ~~The following~~ Table 2 summarizes the results:

	Using Joint Optimization	Using Individual Optimization
Drive Yield (Yd)	93.75	90.625
Drives failed after VZ	4 & 29	4, 6 & 29
Drives recovered	2, 3, 13, 19, 21 & 25	2, 3, 13, 19, 21 & 25
Passed drives failed after VZ	None	6
Drives failed before VZ	2, 3, 4, 12, 19, 21 & 29	2, 3, 4, 12, 19, 21 & 29
	i.e., drive yield before VZ Yd=75%	i.e., d~~D~~rive yield after VZ Yd = 75%

Table 2

~~In addition to disk drives, the present invention is useful with other storage devices such as e.g. tape drives, optical drives, etc.~~ A-It~~T~~hough a manufacturing test case for a two format design is illustrate~~described~~d, the search algorithm can be easily ~~be~~ generalized to a higher number of formats. The design of two formats based on 1D, 2D and 3D ~~joint~~ ~~storage density~~ BPI distributions can easily be generalized to higher order or dimensions by considering more ~~zones~~ than three zones. The ~~design of~~ format designs can be generalized from two to a higher number of formats. The measurement procedure can be generalized to consider more zones as well as off-track measurements such as 747 curves or quality metrics versus error rate measurements to perform a correlation study for the ~~choice of~~ best metric with the leastss ~~potential~~ test time.

Further, the ~~above methods for a~~ per zone variable BPI design can be easily extended to a variable BPI/TPI design ~~as described below~~. The measurement process is extended to ~~further~~ include 747 measurements of all the heads from a pre-selected number of disk drives. To speed up the measurements ~~of raw data~~, instead of 747 measurements, off-track and adjacency margin ~~(~~squeeze measurements~~)~~ of the~~all~~ heads can be performed. Once the 747 ~~raw~~ data of the~~all~~ heads at a pre-selected number of zone ~~locations~~ is determined, for every zone, ~~(~~joint~~)~~ BPI/TPI distributions can be extracted at the ~~given desired~~ target(s) by post-measurement data processing ~~of data~~. The choice of a target is an integral part of the ~~amount of~~ performance gain, such as ~~disk drive~~ yield, due to the per zone variable BPI/TPI designs. Some example ~~choices of~~ target~~(s)~~ are off-track symbol error rate, ~~the variance~~

~~of~~ position error signal variance, and ~~or even~~ a combination of both. After the joint BPI/TPI distributions are extracted and available for the~~all~~ zones, a per zone variable BPI/TPI design can be obtained by solving two-constrained (joint) optimization problems: one that maximizes the ~~drive~~ yield while keeping the same disk drive areal density, and another that maximizes the disk drive areal density while keeping the same ~~drive~~ yield. Once the per zone variable BPI/TPI designs are obtained, ~~a head BPI/TPI allocation and selection criterion, similar to that described herein, can be used such that~~ a pre-selected number of heads are allocated to high and low density BPI and TPI formats, for example, for ~~the case of~~ a two variable BPI/TPI per zone design performed as part of the test process.

The present invention improves storage capacity ~~drive yield and drive capacity~~ (and ~~or~~ consequently areal density at a fixed target BPI) and yield~~,~~ and ~~allows~~ reduc~~es~~ing the target TPI by increasing the average BPI across the stroke per head (depending on the number of formats considered) to meet a desire storage ~~target drive~~ capacity. In particular, ~~due to a maximum deliverable data rate of the ASIC components (e.g., channel, controller and preamp)~~ the BPI at the outer diameter may be limited by the maximum~~minimum~~ deliverable data rate of the ~~mentioned~~ ASIC components. -For example, if the controller 57 has~~is capable of~~ a maximum deliverable data rate of 650 MHz, the preamplifier 21 has~~capable~~ a maximum deliverable data rate of 700 MHz and the read/write channel 51 has~~capable~~ a maximum deliverable data rate of 750 MHz, then the BPI at the outer diameter is limited by the controller 57 at a maximum deliverable data rate of 650 MHz. Thus, a conventional one format BPI profile across the stroke does not achieve the desired storage ~~drive~~ capacity and ~~the desired manufacturing drive~~ yield~~. The target BPI is increased and the BPI profile across the stroke is relaxed, wherein according to~~ whereas the present invention~~, the~~ per zone variable BPI ~~design can be used to design (variable BPI)~~ target formats ~~that~~ meet the desired storage ~~drive~~ capacity at a fixed target TPI while improving the ~~overall drive~~ yield.

~~Referring back to FIG. 2A and FIGs. 3-6, in one embodiment of the present invention, the steps of the example method of the present invention: (1)~~ Tthe data ~~collection/~~measurer~~ment process block~~ 62~~,~~ can be implemented by ~~on~~ a general purpose computer~~ing equipment~~ 61~~, known in the art,~~ and the drive electronics of the disk drive 100. The general purpose computer 61 can be a high end PC, a PC server or a workstation and include programmable simulation software. The drive electronics can include~~including~~ the ~~special purpose electronic circuit (e.g.,~~ logic circuit~~)~~ 49 and the controller ~~on board microprocessor~~ 57. T-he logic ~~(FIG. 1B), configured according to the present invention, wherein the special purpose electronic~~ circuit 49 ~~is configured to~~ performs the measurements and~~,~~ the controller~~on board microprocessor~~ 57 directs the logic ~~special purpose~~ circuit 49~~,~~ and transfers the data to the general purpose computer 61. T~~, (2)~~ the head assignments ~~process~~ can be implemented by~~on~~ the controller ~~on board microprocessor~~ 57 with ~~within the disk drive 100 configured according to the present invention,~~ the ~~wherein a~~ data collection sub-task ~~is related to the head assignment task such that the data collection sub-task is~~ performed by the logic~~special purpose electronic~~ circuit 49 ~~within the disk drive 100~~. The ~~, (3) the steps in each of the~~ post-measurement data -processor~~ing block~~ 64, the format optimizer ~~block~~ 66 and the format generator ~~block~~ 68 can be implemented by~~on~~ the general purpose computer~~ing equipment~~ 61 ~~(e.g., high end PC, PC server or workstation, etc., including programmable simulation software) configured according to the present invention~~.

The present invention has been described in considerable detail with reference to certain preferred versions thereof; however, other versions are possible. _Therefore, the spirit and scope of the appended claims should not be limited to the description of the preferred versions contained herein.

Abstract

A method of defining ~~such~~ a storage format in ~~multiple~~ data storage devices,- each data storage device having multiple ~~a plurality of~~ storage media and ~~a plurality of~~ corresponding ~~data transducer~~ heads, each ~~transducer~~ head for recording on and playback of information from a corresponding storage media~~um~~ in at least one zone, ~~wherein~~ and each zone including~~es a plurality of~~ concentric tracks for recording on and playback of information. The method includes ~~the steps of:~~ selecting a sample of the~~a plurality of said~~ data storage devices,~~;~~ for each selected data storage device~~,~~ measuring a record/playback performance capability of each head at one or more read/write frequencies per zone,~~;~~ ~~based on said performance capability measurements,~~ generating storage density distributions corresponding to ~~at least a number of~~ the heads in the~~said~~ selected data storage devices based on the performance capability measurements,~~;~~ selecting a group of read/write frequencies for the~~said~~ ~~multiple~~ data storage devices with~~,~~ two or more frequencies for each zone~~,~~ based on the~~said~~ storage density distributions,~~;~~ and assigning one of the~~said~~ read/write frequencies to each head based on the performance capability of that head.

SUBSTITUTE SPECIFICATION UNDER 37 C.F.R. 1.125

Per Zone Variable BPI for Improving Storage Device Capacity and Yield

By

Saeed Asgari

Mathew Vea

And

George Iszlai

Field of the Invention

The present invention relates to information storage on a storage media such as a disk in a disk drive.

Background of the Invention

Data storage devices such as disk drives are used in many data processing systems. Typically a disk drive includes a magnetic data disk having disk surfaces with concentric data tracks, and a transducer head paired with each disk surface for reading data from and writing data to the data tracks.

Disk drive storage capacity increases by increasing the data density (or areal density) of the data stored on the disk surfaces. Data density is the linear bit density on the tracks multiplied by the track density across the disk surface. Data density is measured in bits per square inch (BPSI), linear bit density is measured in bits per inch (BPI) and track density is measured in tracks per inch (TPI). As data density increases, the head performance distribution also increases which diminishes disk drive storage capacity and yield.

Conventional disk drives fail to account for the different capabilities of the head and disk surface pairs. Conventionally, each disk surface is formatted to store the same amount of data as every other disk surface. However, each head and disk surface pair has unique data recording capability, such as sensitivity

and accuracy, which depends on the fly height of the head over the disk surface, the magnetic properties of the head and the quality/distribution of the magnetic media for the disk surface. Thus, in conventional disk drives a head and disk surface pair that has a low error rate is formatted to the same BPI and TPI as a head and disk surface pair that has a high error rate.

Conventional disk drive manufacturing applies a single error rate and a single data storage level for the head and disk surface pairs, and scraps disk drives that include a low performing head and disk surface pair that fails to meet the qualifying requirements. This lowers storage capacity due to inefficient use of high performing head and disk surface pairs that can store more data, and lowers yield due to disk drives being scrapped if they include a low performing head and disk surface pair even if they also include a high performing head and disk surface pair.

U.S. Patent Nos. 6,091,559 and 5,596,458 provide different BPI on different disk surfaces, however these approaches do not take into consideration multiple constraints, including head performance across the stroke per disk surface, performance requirements such as throughput and manufacturing requirements such as test time. Instead, disk surface zone frequencies are selected based on a single metric for one head.

There is, therefore, a need for storing data in a disk drive which improves disk drive storage capacity and yield and accounts for head performance variation.

Summary of the Invention

The present invention satisfies this need.

In an embodiment, a variable BPI storage format is a function of zones in data storage devices, such as disk drives, based on head performance variation between different heads in a set of data storage devices.

In another embodiment, a population of disk drives is selected, and head performance measurements are taken for disk surface locations at different frequencies. Head performance distributions obtained from the head performance measurements provide storage formats for the disks by determining different read/write frequencies for the zones, and the heads in each disk drive are assigned to the frequencies.

The head allocations and assignments are per head per zone, taking into consideration head performance variation across the zones. For instance, if a first head performs well at the inner diameter (ID) of the disk but poorly at the outer diameter (OD) of the disk, and a second head has reverse performance, then the first head is assigned a high BPI at the ID and a low BPI at the OD, and the second head is assigned in the opposite fashion. The per zone variable BPI storage format improves storage capacity by taking several manufacturing and customer constraints into consideration. Performance of each head across the stroke of the disk surface, as well as performance variation from one head to another, determines the storage format and the head assignments.

In another embodiment, the head performance and the storage format are determined off-line at development/design time, and then the heads are assigned to the different frequencies at manufacturing time. For example, the storage format for each zone and the number of heads allocated to each data density are preselected at design time, and then the heads are assigned to high/low data density storage formats at manufacturing time.

In another embodiment, a method defines the storage format in data storage devices, with each data storage device having multiple storage media

and corresponding heads, each head for recording on and playback of information from a corresponding storage media in multiple zones, and each zone including concentric tracks for recording on and playback of information. The method includes (1) selecting a sample of the data storage devices, (2) for each selected data storage device, measuring a record/playback performance of each head at one or more read/write frequencies per zone, (3) generating head performance distributions based on the head performance measurements, (4) selecting a group of read/write frequencies for the data storage devices, two or more frequencies for each zone, based on the head performance distributions, and thereafter, during manufacturing, (5) assigning one of the read/write frequencies to each head based on the performance of that head.

Advantageously, the present invention provides consistent performance (both sequential and random throughput) across a population of disk drives, improves storage capacity and yield and reduces test time.

Brief Description of the Drawings

These and other features, aspects and advantages of the present invention will become understood with reference to the following description, appended claims and accompanying figures where:

FIG. 1A shows a disk drive with a data storage format;

FIG. 1B shows drive electronics for the disk drive;

FIG. 1C shows servo tracks and data tracks on a disk surface;

FIG. 1D shows a zone format in the disk drive with N disks, 2N heads and different heads in a zone on different disk surfaces;

FIG. 1E shows another zone format on a disk surface;

FIG. 1F shows zones on a disk surface that each include virtual cylinders;

FIG. 1G shows a data track format for the virtual cylinders in a zone on different disk surfaces with corresponding heads;

FIG. 1H shows a servo track and data track format for a zone on different disk surfaces with corresponding heads in which the number of servo tracks and

data tracks in different virtual cylinders of a zone on different disk surfaces are the same;

FIG. 1I shows another servo and data track format that varies from zone to zone on a disk surface;

FIG. 2A shows a function and flow diagram for generating the format of FIG. 1A;

FIG. 2B shows a graph of playback error measurement for a head at a zone at different recording frequencies;

FIG. 2C shows a histogram of the frequency capabilities of the heads in a set of disk drives at a zone at a fixed target error rate;

FIG. 2D shows a joint BPI distribution;

FIG. 3 shows a flowchart of vertical zoning data collection in FIG. 2A;

FIG. 4 shows a flowchart of vertical zoning and per zone joint BPI distribution extraction in FIG. 2A;

FIG. 5 shows a flowchart of head assignments in FIG. 2A; and

FIG. 6 shows a flowchart of format generation and optimization in FIG. 2A.

Detailed Description of the Invention

Data storage devices used to store data for computer systems include, for example, disk drives, floppy drives, tape drives, optical and magneto-optical drives and compact drives. Although the present invention is illustrated by way of a disk drive, the present invention can be used in other data storage devices and other storage media, including non-magnetic storage media, is as apparent to those of ordinary skill in the art and without deviating from the scope of the present invention.

FIGs. 1A-1C show a hard disk drive 100 diagrammatically depicted for storing user data and/or operating instructions for a host computer 54. The disk drive 100 includes an electro-mechanical head-disk assembly 10 that includes one or more rotating data storage disks 12 mounted in a stacked, spaced-apart

relationship upon a spindle 13 rotated by a spindle motor 14 at a predetermined angular velocity.

Each disk 12 includes at least one disk surface 23, and usually two disk surfaces 23 on opposing sides. Each disk surface 23 has associated magnetic media for recording data. The spindle motor 14 rotates the spindle 13 to move the disks 12 past the magnetic transducer heads 16 suspended by the suspension arms 17 over each disk surface 23. Generally, each head 16 is attached to a suspension arm 17 by a head gimbal assembly (not shown) that enables the head 16 to swivel to conform to a disk surface 23. The suspension arms 17 extend radially from a rotary voice coil motor 20. The voice coil motor 20 rotates the suspension arms 17 and thereby positions the heads 16 over the appropriate areas of the disk surfaces 23 in order to read from or write to the disk surfaces 23. Because the disks 12 rotate at relatively high speed, the heads 16 ride over the disk surfaces 23 on a cushion of air (air bearing).

Each head 16 includes a read element (not shown) for reading data from a disk surface 23 and a write element (not shown) for writing data to a disk surface 23. Most preferably, the read element is a magneto-resistive or giant magneto-resistive sensor and the write element is inductive and has a write width which is wider than a read width of the read element.

Each disk surface 23 is divided into concentric circular data tracks 30 that each have individually addressable data sectors 35 in which user data is stored in the form of magnetic bits. The data sectors 35 are separated by narrow embedded servo sectors 25 arranged in radially extending servo spokes. The servo sectors 25 include a series of phase-coherent digital fields followed by a series of constant frequency servo bursts. The servo bursts are radially offset and circumferentially sequential, and are provided in sufficient numbers that fractional amplitude read signals generated by the head 16 from portions of at least two servo bursts passing under the head 16 enable the controller 57 to

determine and maintain proper position of the head 16 relative to a data track 30. A servo burst pattern for use with a head that includes a magneto-resistive read element and an inductive write element is described by commonly assigned U.S. Patent No. 5,587,850 entitled "Data Track Pattern Including Embedded Servo Sectors for Magneto-Resistive Read/Inductive Write Head Structure for a Disk Drive" which is incorporated herein by reference.

The controller 57 controls the heads 16 to read from and write to the disk surfaces 23. The controller 57 preferably is an application specific integrated circuit chip (ASIC) which is connected by a printed circuit board 50 to other ASICs, such as a read/write channel 51, a motor driver 53 and a cache buffer 55. The controller 57 preferably includes an interface 59 which connects to the host computer 54 via a known bus 52 such as an ATA or SCSI bus.

The controller 57 executes embedded or system software including programming code that monitors and operates the disk drive 100. During a read or write operation, the host computer 54 determines the address where the data is located in the disk drive 100. The address specifies the head 16, the data track 30 and the data sector 35. This data is transferred to the controller 57 which maps the address to the physical location in the disk drive 100, and in response to reading the servo information in the servo sectors 25, operates the voice coil motor 20 to position the head 16 over the corresponding data track 30. As the disk surface 23 rotates, the head 16 reads the servo information embedded in each servo sector 25 and also reads an address of each data sector 35 in the data track 30.

During a read operation, when the identified data sector 35 appears under the head 16, the entire contents of the data sector 35 containing the desired data is read. In reading data from the disk surface 23, the head 16 senses a variation in an electrical current flowing through the read element when it passes over an area of flux reversals on the disk surface 23. The flux reversals are transformed

into recovered data by the read/write channel 51 in accordance with a channel algorithm such as partial response, maximum likelihood (PRML). The recovered data is then read into the cache buffer 55 where it is transferred to the host computer 54. The read/write channel 51 most preferably includes a quality monitor which measures the quality of recovered data and provides an indication of the data error rate. One channel implementation which employs channel error metrics is described in commonly assigned U.S. Patent No. 5,521,945 entitled "Reduced Complexity EPR4 Post-Processor for Sampled Data Detection" which is incorporated herein by reference. The present invention uses the indication of recovered data error to select linear bit density, track density and/or error correction codes.

During a write operation, the host computer 54 remembers the address for each file on the disk surface 23 and which data sectors 35 are available for new data. The controller 57 operates the voice coil motor 20 in response to the servo information read back from the servo sectors 25 to position the head 16, settles the head 16 into a writing position, and waits for the appropriate data sector 35 to rotate under the head 16 to write the data. To write data on the disk surface 23, an electrical current is passed through a write coil in the inductive write element of the head 16 to create a magnetic field across a magnetic gap in a pair of write poles that magnetizes the disk surface 23 under the head 16. When the data track 30 is full, the controller 57 moves the head 16 to the next available data track 30 with sufficient contiguous space for writing data. If still more track capacity is required, another head 16 is used to write data to a data sector 35 of another data track 30 on another disk surface 23.

The present invention increases the storage capacity and yield of data storage devices, such as the disk drive 100, having magnetic media surfaces, such as the disk surfaces 23.

Vertical Zoning

In every disk drive, there is a distribution associated with the head and disk surface pair performance. The present invention takes advantage of that distribution to determine different linear bit density (BPI) recording frequency assignments for the heads, and optionally track allocation.

A set of disk drives is selected, and head performance measurements are taken for each selected disk surface location in the disk drives at different frequencies. Empirical frequency capability histograms are extracted at a target performance metric from the measurement data. Head performance distributions (such as joint BPI distributions) are estimated from the histograms and fed into a format optimizer to obtain and design vertically zoned frequency format profiles across the stroke and the disk surface as well as the optimal number of head allocations to the frequencies. Once the frequency format profiles and the optimal number of head allocations are determined, during a test process, every head at every zone is assigned to one of the frequencies based on the head's performance.

FIG. 1A shows a storage format for the disk drive 100. Each disk surface 23 includes zones 60 that extend from one radius of the disk 12 to another radius of the disk 12, and the format of the zones 60 on each disk surface 23 is the same. The variable BPI storage format is a function of the zones 60 on each disk surface 23 based two data recording formats -- high data density and low data density -- that use (1) head performance variation from one head 16 to the next head 16 in the disk drive 100, and (2) the performance variation of a given head 16 across the stroke of a disk surface 23.

The disk drive 100 includes the disks 12 depicted as disks 1 to N, the heads 16 depicted as heads 1 to 2N, and the disk surfaces 23 depicted as disk surfaces 1 to 2N. Each disk 12 includes two opposing disk surfaces 23, and each head 16 is associated with one of the disk surfaces 23. For instance, head 1 is associated with disk surface 1 of disk 1, head 2 is associated with disk

surface 2 of disk 1, head 3 is associated with disk surface 3 of disk 2, and head 2N is associated with disk surface 2N of disk N.

Each disk surface 23 includes the zones 60 depicted as zones 1 to M across its stroke, with zone 1 at the ID and zone M at the OD. The radial boundaries on zone 1 of disk surface 1 of disk 1 are the same as the radial boundaries of zone 1 on disk surface 2 of disk 1, and so on. Similarly, the radial boundaries of zone M on disk surface 1 of disk 1 are the same as the radial boundaries of zone M on disk surface 2 of disk 1, and so on. However, different zones 60 across the stroke on each disk surface 23 need not necessarily have the same number of data tracks 30 or TPI. For example, zone 1 on disk surface 1 of disk 1 has the same number of data tracks 30 and the same radial boundaries as zone 1 on disk surface 1 of disk N, and zone M on disk surface 1 of disk 1 has the same number of data tracks 30 and the same radial boundaries as zone M on disk surface 1 of disk N. However, the number of data tracks 30 in zones 1 and M can be different.

Each disk surface 23 also includes virtual cylinders 39 depicted as virtual cylinders 1 to n. Each zone 60 includes multiple virtual cylinders 39, and each virtual cylinder 39 includes multiple data tracks 30 on each disk surface 23. Further, within a virtual cylinder 39, different heads 16 may read and write at different frequencies (variable BPI) to provide vertical zoning.

FIG. 1C shows the data tracks 30 and the servo tracks 37 on the disk surface 23. The data tracks 30 include the data sectors 35, and the servo tracks 37 include the servo sectors 25. Five servo tracks 37 depicted as servo tracks Sa, Sb, Sc, Sd and Se are shown in relation to three data tracks 30 depicted as data tracks Tk1, Tk2 and Tk3.

The servo tracks 37 are written on the disk surface 23 during manufacturing at a servo track density that is about 150% of the maximum data

track density. The servo track density is determined by the maximum read width and the minimum write width of a population of the heads 16. After writing the servo tracks 37 at the servo track pitch, the data tracks 30 can be written at any radial position between the servo tracks 37. The data track density (TPI) can be selected from predetermined levels or can be based on the location of a data sector 35. Additional tests can be performed to determine the optimum data track density of the disk surface 23. Each servo track 37 comprises radially similarly situated servo sectors 25 in the servo spokes. For example, the servo track Se contains servo sectors 25 at essentially the same radial distance from the center of the disk 12, the servo track Sd contains servo sectors 25 at essentially the same radial distance from the center of the disk 12, etc.

FIGs. 1D to 1I show vertical zone formats in which different heads 16 on different disk surfaces 23 may read/write at different linear frequencies (variable BPI) on the data tracks 30 within a virtual cylinder 39.

FIG. 1D shows a zone 60 format of the disk drive 100 with N disks 12, 2N heads 16 and different heads 16 in zone 1 on different disks 12. FIG. 1E shows another zone 60 format on the disk surface 23. FIG. 1F shows each zone 60 on the disk surface 23 includes multiple virtual cylinders 39. Zone 1 includes virtual cylinders 1 to j, and zone M includes virtual cylinders 1 to i. The radial boundaries of the zones 60 are shown as dark circles, and the radial boundaries of the virtual cylinders 39 are shown as light circles. FIG. 1G shows a data track 30 format for the virtual cylinders 39 in a zone 60 on different disk surfaces 23 with corresponding heads 16. FIG. 1H shows a data track 30 and servo track 37 format for a zone 60 on different disk surfaces 23 with corresponding heads 16 in which the number of data tracks 30 and servo tracks 37 in different virtual cylinders 39 of a zone 60 on different disk surfaces 23 is the same. FIG. 1I shows a data track 30 and servo track 37 format for zones 60 on the disk surface 23 with a corresponding head 16 in which the data track 30 and servo track 37 format varies from zone 1 to zone M.

Format Optimization

The vertical zoning includes designing, optimizing and selecting two or more recording frequency profiles per zone for a sample number of disk drives off-line during the disk drive development/design phase. Then, for a population of disk drives, in each disk drive, each head is assigned to one of the predetermined frequencies for a given zone during the disk drive manufacturing phase. A predetermined read/write frequency (BPI) is assigned to each head based on a known number of head allocations and the head's performance. A head assigned to a high frequency records more bits on a track, and a head assigned to a low frequency records less bits on a track.

Performance testing of the head and disk surface pairs occurs after full read/write and servo calibration and optimization of the disk drive. If the tested performance of head 1 at zone 1 on disk surface 1 of disk 1 at a given frequency is better than a target performance metric, then head 1 is considered strong since it is capable of storing more information than originally accounted for. Thus, the recording frequency can be increased at zone 1 on disk surface 1 of disk 1 for head 1 yet the performance does not fall below the target performance metric. If the tested performance of head 2 at zone 1 on disk surface 2 of disk 1 at the same frequency is worse than a target performance metric, then head 2 is considered weak but can be compensated for by relaxing the frequency at which head 2 operates to ensure the target performance metric is met. Performing the above trade-off between the heads for all the zones provides frequency profiles across the stroke that are vertically zoned frequency format profiles without loss of storage capacity.

Advantageously, by compensating for head 2, rather than failing the disk drive due to head 2, the vertical zoning improves yield. Furthermore, the format optimizer uses the head performance (read/write frequency capability) distributions at every zone and a target performance metric to design a group of

read/write frequency format profiles for strong and weak heads within a given disk drive. The format optimizer also determines the optimal number of strong versus weak heads.

The format optimizer does not determine which specific head is at the high or low frequency, but does provide a breakdown of the number of heads at the high frequency and the number of heads at the low frequency. The breakdown is fixed, performed off-line, and used during the head assignments. Then, in the head assignments during a manufacturing test, out of 2N heads in a disk drive with N disks, the number of heads assigned to each predetermined frequency is determined.

The heads within a set of disk drives are allocated to the predetermined group of read/write frequencies as part of the optimization process to meet the storage capacity and yield requirements for the disk drives. The allocation process allocates a number of the heads in a disk drive to the predetermined frequencies, however the specific assignment of a particular head to a particular frequency is performed later during the assignment process. For example, in a two read/write frequency design (high frequency and low frequency) for a set of disk drives each with eight heads, in each disk drive for zone 1 on all the disk surfaces, any five of the eight heads are allocated to the high frequency and any three of the eight heads are allocated to the low frequency based on the performance measurements of the heads in the set of the disk drives. Thereafter, the specific assignment of each particular head to a particular predetermined frequency is performed. For example, in a first disk drive heads 2, 5, 6, 7, 8 are assigned to the high frequency and heads 1, 3, 4 are assigned to the low frequency, whereas in a second disk drive heads 1, 4, 5, 6, 7 are assigned to the high frequency and heads 2, 3, 8 are assigned to the low frequency. The specific head assignments depend on the specific capability of the heads in each disk drive.

The optimal number of heads per frequency is determined at the same time that the group of read/write frequencies are selected by the format optimizer by solving a joint constrained optimization problem. For example, in a disk drive with eight heads and a high frequency and a low frequency that are each a different ratio of a reference frequency, in each vertical zone, allocating two heads to the high frequency and six heads to the low frequency provides a specific storage capacity. Changing the frequency ratios and the number of heads allocated to each frequency provides a different storage capacity. Thus, the disk drive storage capacity is a function of the number of heads multiplied by the frequency allocated to each head per zone. For example, if a nominal disk surface data storage is 1 unit, and if the high frequency = 4/3 x the reference frequency and the low frequency = 2/3 x the reference frequency, then one head can be at the high frequency for every one head at the low frequency to maintain the average disk surface data storage at 1 unit.

The head performance distributions represent percentages of the heads in the disk drives than can operate at different frequencies. For example, the head performance distribution is a BPI distribution that represents the head frequency capability at a target performance metric. Using the head performance distributions (the head read/write frequency capability distributions at the target performance metric for every zone), the number of heads, the format of the virtual cylinders and the desired storage capacity, the format optimizer determines the frequency for each virtual cylinder in each zone and the number of heads in each disk drive allocated to each frequency to achieve the desired storage capacity. Thereafter, in the assignment process as part of testing each disk drive, each head in a population of disk drives is assigned to one of the predetermined frequencies based on the allocation criteria and the specific head performance. For example, in a disk drive with four heads, the format optimizer considers three heads at the high frequency and one head at the low frequency, then two heads at the high frequency and two heads at the low frequency, and then one head at the high frequency and three heads at the low frequency.

Thus, the format optimizer uses the head performance distributions to determine the storage capacity and yield.

In one version of the optimization process, the yield is maximized while meeting a constraint on storage capacity. In another version, the storage capacity is maximized while meeting a constraint on yield. In the former case, the format optimizer uses a format where the maximum number of disk drives qualify and the fewest number of disk drives fail to reach the required storage capacity. For example, in a disk drive with four heads and a nominal disk surface data storage of 1 unit, allocating two heads to the high frequency and two heads to the low frequency provides a nominal data storage of 4 units. In the later case, the format optimizer uses a format where the maximum number of disk drives reach the required storage capacity and the fewest number of disk drives fail to qualify. For example, in a disk drive with four heads and a nominal disk surface data storage of 1 unit, allocating three heads to the high frequency and one head to the low frequency provides a high data storage of 4.66 units.

Thus, the vertical zoning for variable BPI includes an off-line predetermined per zone format design based on disk drive data collection and head performance distribution extraction. In one version, a fixed predetermined zone boundary format is used to design multiple frequency BPI formats based on representative or actual joint BPI distributions at one or more desired target performance metrics (such as off-track symbol error rate) and the joint BPI distributions are extracted from a finite preselected set of disk drives.

The collected data is used to extract the joint BPI distributions for the heads at every preselected zone, and the per zone design of high and low data density formats for the heads is performed off-line. The format optimizer solves a constrained joint optimization off-line to obtain the format designs using well-known constrained optimization routines. Using joint BPI distributions allows consideration of potential correlation of BPI capability of the heads across the

stroke as well as the individual contribution of each head to the storage capacity and yield.

The off-line format design allows the format optimizer to consider additional constraints. For example, as more information is obtained in quantifying the thermal stability constraints of the disks (which in turn places an upper bound on linear bit density for the heads) the off-line format design does not exceed these constraints. Likewise, if there are data rate constraints in either the write process or the ASICs, such constraints may be cast within the joint constrained format optimizer to ensure the constraints are not exceeded.

Data Measurement

A measurement procedure is used to collect data from which one-dimensional (1D), two-dimensional (2D) and three-dimensional (3D) BPI distributions at a desired read/write target error rate (or any other metric) can be extracted. Data is collected based on head capability measurements taken at different radial positions on the disk. The distributions represent the capability of each head at different radial positions. For example, several disk drives which collectively include 1000 heads are selected for measurement, and record/playback error rate measurements of the 1000 heads from zone 1 to zone 24 of the disk surfaces at different frequencies are obtained. Thereafter, in post-measurement data processing (a) the BPI capability of each head at a fixed target performance metric at zone 1 is determined to obtain a 1D BPI distribution, (b) the BPI capability of each head at a fixed target performance metric at zones 1 and 5 is determined to obtain a 2D BPI distribution, and (c) the BPI capability of each head at a fixed target performance metric at zones 1, 5 and 20 is determined to obtain a 3D BPI distribution.

The BPI distributions are then passed to the format optimizer to solve three constrained optimization problems to provide head frequency per zone allocations. The three constrained optimization problems (1) maximize the yield

while preserving the storage capacity, (2) maximize the storage capacity while preserving the yield, and (3) maximize the yield while ensuring a target storage capacity is met at a fixed target TPI. Customer related or ASIC data rate constraints are also considered. The format optimizer can solve any of these three problems, and one problem can take priority over another depending on the process phase. For example, at an early development phase where the disk drive components are not mature, meeting the storage capacity may be a challenge. In that phase, the format optimizer can design the variable BPI format profiles by solving the second problem. Then, as the disk drive components mature, meeting the storage capacity becomes easier and meeting the yield becomes more important, the first problem may be solved. Thereafter, as part of a test process, an assignment algorithm ensures the appropriate head assignments to the predetermined high and low data density formats per head and per zone or across the stroke based on the head allocation breakdown of the format optimizer.

The yield is improved while meeting the target storage capacity by allowing a frequency format with high and low frequencies and a predetermined number of high and low performing head allocations. Using realistic constraints such as ASIC data rate limitations, the same fixed target TPI is maintained by increasing the average target BPI across the stroke to achieve the target storage capacity. As such, head performance variation from one head to the next head in the disk drive and for the head across the stroke of the disk surface is used to increase the storage capacity while preserving the yield. For example, the vertical zoning format uses several design constraints to improve yield using a variable high and low BPI design with a fixed predetermined number of head allocations as a function of the zones while meeting the target storage capacity at a fixed target TPI. The head performance variation or correlation across the stroke is also used.

Further, the difference in data storage of two or mores zones on a disk surface is considered as it affects storage capacity. The storage capacity is defined as a weighted combination of the zone capacities across the stroke on each disk surface in the disk drive. A correlation in the head performance statistics is extracted from one head to another head, and for every head considered in a set of disk drives across the stroke on each disk surface.

The joint constrained optimization determines a per zone target high and low data density format. The optimization takes into account constraints including customer related requirements such as minimum logical block count, monotonic data rate, and maximum data rate at the outer zones which can be formulated into additional constraints.

Example Implementation

FIG. 2A shows a function and flow diagram for generating the optimal data density format shown in FIG. 1A. The function and flow diagram includes a data measurer 62, a post-measurement data processor 64, a format optimizer 66 and a format generator 68.

Data Measurer

The data measurer 62 takes data measurements for every zone at a finite number of frequency samples.

The data measurer 62 implements a measurement procedure that includes the steps of:

(1) Create several different predetermined linear bit density format profiles including a profile of different frequencies per zone across the stroke, such as a first profile including high frequency 1 for zone 1, high frequency 2 for zone 2 . . . high frequency M for zone M, and a second profile including low frequency 1 for zone 1, low frequency 2 for zone 2 . . . low frequency M for zone M to be loaded on a representative number of disk drives selected for the

measurement process (or if possible on all the available disk drives for that build);

(2) Load a frequency format profile;

(3) Perform read/write and servo optimization and calibration;

(4) Take head performance measurements including off-track mean square error or quality metric and/or symbol error rate at preselected frequencies for preferably all available zones and save the data; and

(5) Repeat steps 2-4 for the remaining frequency format profiles.

The above steps are performed for the selected disk drives in the measurement process.

Thus, in the disk drive 100, the data is recorded on a data sector 35 of the disk surface 23 at the selected data density by positioning the head 16 abutting the data sector 35 and sending the appropriate write signals to the head 16. Typically, a sample of data is recorded on the disk surface 23 such that a significant number of errors are detected (such as ten errors per error rate measurement) to obtain a statistically representative sampling of the error rate for the data sector 35. Thereafter, the recorded data is read by the head 16 and stored by the host computer 54 for evaluation. An error rate of the recorded data is measured or compiled by comparing the written data with the read data, element-by-element. The error rate can be determined using a bit error measurement in which a bit of data read from the disk surface 23 is compared with the correct bit, a bit steam measurement in which a bit stream of data read from the disk surface 23 is compared with a correct bit stream, or a mean square error metric measurement in which a waveform read from the disk surface 23 is compared with an ideal waveform to provide an error signal that is squared and summed to form the error metric.

In this description, a component distribution is defined as a random variation (tolerance) of a prespecified target nominal component parameter such

as a head write/read width, and a distribution is defined as a probability distribution function. During the early product development, when the head performance distributions are wide and unreliable, data from a mature set of disk drives is used for extracting reference joint BPI distributions at a target performance metric such as on-track symbol error rate, off-track symbol error rate, on-track mean square error or off-track mean square error. Later, when the head performance variation from one phase to the next in the distribution is expected to be minimal, new sets of measurement data are collected using a selected population of disk drives at their more mature stages.

Thus, a number of BPI formats including the nominal target format are selected. Then, on-track or off-track symbol error rate or mean square error measurements are taken at different preselected locations of the disk surfaces, such as the outer, middle and inner zones. The performance measurements can be limited to these three zones to reduce the measurement time. However, preferably the performance measurements over multiple zones and other measurements such as off-track 747 can be performed. The nominal formats are generated from the data.

Two or more different linear bit density format profiles can be loaded at a time. In one example, two variable BPI format per zone design (high and low data density format profiles) can be created for measurement data collection during every build. In this way, more statistical data can be collected from more disk drives, however there will be only two frequency samples per zone available for post-measurement data processing.

Post-Measurement Data Processor

Post-measurement data processor 64 uses the available performance metric measurements to calculate each head's frequency performance, for instance as kilo flux per inch (kFCI) or kilo bits per inch (kBPI), at a target

performance metric. The performance of every head at every zone is determined as a function of the read/write frequency profiles used for the measurements.

For example, if six different frequency profiles are used, then for every head per zone, the data measurer 62 provides measured data as a function of six frequency samples at a target performance metric. In the post-measurement data processor 64, the measured data is sorted and the performance of every head at every zone at the six frequency samples is extracted to generate frequency capability histograms at a target performance metric.

FIG. 2B shows a graph of playback error measurement for a head at a zone at different recording frequencies. The curve shows head performance as a function of frequency (BPI). The x-axis is the read/write frequency in kBPI at the outer diameter, and the y-axis is the on-track symbol error rate on a log scale. Each frequency sample 70 is depicted as "+", each curve fit point 72 is depicted as "o" and each projected frequency 74 is depicted as "◊".

In the illustration, head 1 at zone 1 in disk drive 3 is measured at six frequency samples. The curve is generated using a least square polynomial fit to the six frequency samples. The projected frequency (BPI) for a target on-track symbol error rate is extracted from the curve by interpolation or extrapolation. For example, if the target on-track symbol error rate is 10^{-8} then the projected frequency is determined by interpolation, whereas if the target on-track symbol error rate is 10^{-6} then the projected frequency is determined by extrapolation. The on-track symbol error rate varies as a function of frequency and increases as the frequency increases.

The nominal kBPI (before vertical zoning) and the kBPI gain relative to the nominal kBPI are also shown. Head 1 can be classified as a strong head because there is reasonably significant margin before its on-track symbol error rate of –9.1 (log) at a nominal frequency/kBPI of ~ 188 can be changed to a

projected on-track symbol error rate of –6.22 at a frequency/kBPI of ~ 217. Hence, there is a total kBPI gain of ~ 29, allowing the nominal frequency to increase by 15% while meeting the target on-track symbol error rate performance metric of $6x10^{-7}$. Thus, head 1 of disk drive 3 has a frequency capability of about 217 at the target on-track symbol error rate of $6x10^{-7}$, which provides a sample for the generation of a histogram.

FIG. 2C shows a histogram of the frequency capabilities of the heads in a set of disk drives at a zone at a target performance metric. The histogram 76 is constructed using the projected frequencies determined in FIG. 2B for the heads in the selected disk drives reading from zone 1 at the target on-track symbol error rate of $6x10^{-7}$. The x-axis is the projected frequency capability at the outer diameter, and the y-axis is the number of heads. The histogram is extracted and empirical, has a normal distribution fit and has a width that corresponds to the head performance variation.

Additional histograms are constructed for the remaining zones based on the frequency capabilities determined from the graphs based on the performance measurements taken at the remaining zones so that every available head considered in the disk drives under measurement has BPI histograms at a target performance metric per zone.

Thus, performance measurements are provided for each head at each zone in the selected disk drives, the graphs are generated for each head at each zone, the frequency capabilities for each head at each zone are determined for a target performance metric, and the histograms are constructed for each head at each zone for the target performance metric. Likewise, if a histogram of head BPI capability at a target performance metric of a zone (such as an intermediate zone) is not available then the histogram for that zone can be constructed by interpolation or extrapolation. The histograms can be used to estimate a BPI distribution.

FIG. 2D shows a joint BPI distribution calculated from the histograms of the heads in the measured disk drives at a target performance metric. The joint BPI distribution is a 2D distribution based on the histograms in FIG. 2C at the target on-track symbol error rate of 6×10^{-7}. The x-axis is the BPI capability of the heads at the middle diameter (MD) of the disks, the y-axis is the BPI capability of the heads at the outer diameter (OD) of the disks, and the z-axis is the calculated number of heads divided by the total number of heads. The joint BPI distribution provides an estimate of the probability that the heads meet the target performance metric at the MD and the OD.

The joint BPI distribution may predict, for example, that 10% of the heads in the measured disk drives can operate at a high frequency of 1.5 x the reference frequency, 50% of the heads can operate at a high frequency of 1.25 x the reference frequency, 90% of the heads can operate at the reference frequency, and 99.9% of the heads can operate at a low frequency of 0.75 x the reference frequency.

For example, the linear bit density sensitivity of every head at zone K (where K ranges from 1 to M) at the six frequency samples is determined. If frequency 1K, frequency 2K . . . frequency 6K are the frequency samples at zone K, every head is positioned on the same track in zone K and the record/playback performance of each head is measured at every frequency sample using a target performance metric.

The BPI distributions can be calculated at the target performance metric as ID, 2D or 3D distributions that are marginal, individual or per zone distributions, respectively. The format optimizer uses the estimated frequency capability BPI distributions for every zone at the target performance metric to determine the storage capacity and yield.

Format Optimizer

The format optimizer 66 provides variable BPI optimization. The format optimizer 66 solves three constrained optimization problems in response to various inputs. The first problem maximizes the yield while preserving the storage capacity, the second problem maximizes the storage capacity while preserving the yield, and the third problem maximizes the yield while reducing the track density and meeting the storage capacity. The inputs include the number of different read/write frequencies (frequency profiles or formats), the number of heads in each disk drive, the BPI distributions, and the nominal storage capacity. The BPI distributions indicate the frequency capability distribution of the heads at a target performance metric.

The format optimizer 66 simultaneously searches through a continuous range of all possible frequency capabilities to maximize the yield such that the nominal storage capacity is met. The format optimizer 66 can also perform the same operation with the storage capacity and the yield interchanged.

The format optimizer 66 can optimize high and low data density as a function of the zones. For example, in a disk drive with eight heads, the possibilities are one head at high data density and seven heads at low data density, two heads at high data density and six heads at low data density, three heads at high data density and five heads at low data density, four heads at high data density and four heads at low data density, one head at low data density and seven heads at high data density, two heads at low data density and six heads at high data density, and three heads at low data density and five heads at high data density. The format optimizer 66 considers all the combinatorial possibilities, in each case solves a constrained optimization problem and chooses the optimal solution among the possibilities. Alternatively, the format optimizer 66 can reach the optimal solution more directly by non-linear mixed-integer programming.

Therefore, once the 1D, 2D and 3D BPI distributions at a target performance metric are passed to the format optimizer 66, the format optimizer 66 solves two problems, (1) maximizing or improving the yield due to the target performance metric while meeting the desired nominal storage capacity, and (2) maximizing the storage capacity while meeting the desired nominal yield.

The format optimizer 66 mathematically casts these two problems as constrained optimization problems and solves them using well-known optimization techniques such as a line search algorithm. The constrained optimization problems can also be cast as non-linear mixed-integer programming and solved using existing optimization methods. Example constraints to be considered, and cast mathematically within the format optimizer 66, include not exceeding a certain frequency at the outer diameter due to ASIC data rate limitations or at the inner diameter due to head/disk limitations. Furthermore, closed form equations are derived and used in the format optimizer 66 to estimate the storage capacity and yield. The format generator 66 also considers possible overhead such as adding redundant bits due to error correction coding or gray coding.

The format optimizer 66 also uses information from the format generator 68 such as the calculated format efficiency per zone (defined in percentages as the amount of user data in blocks that can fit in all tracks in a zone), or the number of tracks per zone, to achieve a very close estimate of the storage capacity determined by the format generator 68. Then, the format optimizer 66 calculates optimal linear bit density format profiles as well as the optimal number of heads allocated to each vertically zoned format profile.

For example, histograms are extracted and the corresponding BPI distributions are estimated for different zones at the target on-track symbol error rate of $6\text{x}10^{-7}$. A format design is provided for a disk drive with four heads and two frequencies to optimize yield while meeting storage capacity.

For example, the format optimizer 66 uses the 1D, 2D and 3D BPI distributions at the target performance metric to jointly optimize for vertically zoned frequency format profiles and the corresponding number of head allocations three zones at a time. An advantage of considering three zones instead of one zone, and thus joint optimization instead of individual optimization, is that the joint optimization allows the frequency profiles to be optimized across the stroke on each disk surface. Therefore, joint optimization exploits the potential correlation in performance from one zone to another zone as well as their individual and weighted contribution to the storage capacity. Joint optimization is preferable for a high/low data density format across the stroke for either improving the yield while keeping the same storage capacity or improving the storage capacity while preserving the yield.

The format optimizer 66 generates the target high/low BPI formats per zone, the optimal number of head allocations per format, and an estimate of the storage capacity and yield. The accuracy of the estimates can be sensitive to the underlying BPI distributions at the target performance metric. Further, the target high/low BPI formats can be sensitive to the variance of the BPI distributions. And, the variance of the BPI distributions can be sensitive to the absolute value of the target performance metric and the type of target performance metric. In addition, the target high/low BPI formats are designed three zones at a time and the yield improvement while preserving the storage capacity is based on the profile of the target nominal formats. The format optimizer 66 also allows for smoothing the target variable BPI format designs. The format generator 68 determines the number of tracks per zone, the number of blocks per track, the radius at each zone, as well as block and track format efficiency. This information is saved in output files for use with the format optimizer 66. The format optimizer 66 then saves the target high/low BPI formats per zone that it generates in two separate files that can be loaded into the format generator 68.

Once the target format profiles are calculated, if they are non-smooth across the stroke, optionally a smoothing process is applied. The format profiles are then loaded into the format generator 68 to create vertically zoned formats and configuration pages. The formats and configuration pages are used by the disk drive firmware to create binary files to be loaded into the reserve image of the disk drives as part of the file system. In this fashion, the design and implementation of the format profiles as well as the number of optimal head allocations are performed off-line and are predetermined for every disk drive configuration.

For example, in a disk drive with four heads and four disk surfaces on two disks, the format optimizer 66 designs vertically zoned high and low frequency profiles. Every disk surface is uniformly partitioned into three zones across the stroke, at a track density with a fixed number of tracks per zone, vertically aligned from one disk surface to another. The nominal disk surface data storage before the vertical zoning can be approximated by the sum over all the zones of the nominal tracks per zone multiplied by the nominal BPI per track multiplied by the format efficiency per zone. Format efficiency per zone is the percentage of the user data that is effectively stored per zone. The nominal storage capacity is the nominal disk surface data storage multiplied by the total number of disk surfaces (or heads). The nominal number of tracks per zone and the format efficiency per zone can be generated by the format generator 68.

Performing vertical zoning to improve the yield without losing storage capacity finds the best frequency per zone and per head such that the disk drive meets performance and storage capacity requirements. If a disk drive with four heads fails due to the performance of head 1 at zone 1, but the performance of another head/zone pair, such as head 1 at zone 2 or head 3 at zone 1, is significantly better, passing the tests with reasonable margins, then a higher than nominal frequency at zone 1 or zone 2 is designed for the strong heads and the frequency at zone 1 for the weak head is lowered. This trade-off obtains a

vertically zoned design of variable frequencies per zone such that the storage capacity is preserved. In addition, the number of heads per zone allocated to high or low data density is determined. Thus, the storage capacity can be approximated by the sum over all the zones of the number of strong heads multiplied by the high frequency data storage per zone multiplied by the format efficiency per zone plus the sum over all the zones of the number of weak heads multiplied by the low frequency data storage per zone multiplied by the format efficiency per zone.

For example, the format optimizer 66 is provided with joint BPI distributions at the target performance metric. Then, for every combinatorial possibility of head allocation to high or low frequency, the format optimizer 66 searches through a continuous range of possible frequencies by considering every zone independently using the marginal distributions and by the combination of zones using the joint BPI distributions to maximize the yield calculated using a closed form equation, such that the storage capacity after the vertical zoning is applied is essentially the same as the nominal storage capacity. Further, the optimal high and low frequency profiles for every combination of head allocations is compared and the one that results in the highest yield is chosen and passed to the format generator 68 for the generation of vertically zoned configuration pages to be used by the disk drive firmware.

The disk surfaces can be partitioned into more than three zones. To reduce computational complexity and time, if the selected/designed number of zones per disk surface is more than three, the format optimizer 66 can generate high and low frequency profiles three zones at a time and smooth the profile after post-processing. Another approach includes embedding the smoothing operator in the design and extending the joint optimization to all the zones to consider the impact of smoothing to yield calculation as part of the design rather than the later stages.

In the disk drive with four heads, the yield is maximized while preserving the nominal storage capacity. To determine the number of head allocations, the format optimizer 66 begins with one weak head and three strong heads per zone. The format optimizer 66 searches through a continuous range of possible frequency capabilities per zone, as well as two and three zones at a time, by considering the 1D, 2D and 3D BPI distributions that result in the best calculated yield such that a minimum nominal storage capacity can be obtained. Next, the format optimizer 66 uses two weak heads and two strong heads and repeats solving the constrained optimization problem. This process is continued until all the combinatorial possibilities are considered. Finally, the format optimizer 66 chooses the solution that results in the best yield and provides the target high and low optimal data density format profiles and the associated number of high and low head allocations to the format generator 68. The format generator 68 then generates vertically zoned format files and configuration pages to be used by the disk drive firmware.

Format Generator

The format generator 68 generally performs three functions. First, the format generator 68 uses target formats/frequencies (or linear densities/BPI) for each zone and calculates the data storage of each zone and thus the storage capacity of the disk drive. Second, the format generator 68 calculates the format efficiency (the percent of the disk surface that is occupied by user data) for each zone. Third, the format generator 68 generates configuration pages. The configuration pages contain per-drive, per-zone, and per-head-per-zone parameters that are programmed into the disk drive electronics such as the preamplifier 21, the read/write channel 51 and the controller 57. The parameters are ordered such that the disk drive firmware selects the correct set of parameters to be programmed into each of the components for the particular head and zone that is being written to or read from at the time.

The format generator 68 calculates the frequency and the data storage of each zone taking into consideration limitations in the programmability and the capability of the disk drive components. For example, the heads 16 have varying down-track separation between the read and write elements, the preamplifier 21 has a minimum and maximum delay in turning on the write current, the read/write channel 51 synthesizer frequencies are limited to discrete frequencies, the motor driver 53 can keep the spindle motor 14 within a finite precision of the nominal rotational speed, the controller 57 has specific latencies in generating commands to the preamplifier 21 and the read/write channel 51 often with a finite uncertainty as to the exact timing of these commands, and a reference crystal (not shown) has finite accuracy and stability over temperature.

The format generator 68 can be fully automated, or can be directed by a human operator. In the absence of input from the format optimizer 66, the target per-zone BPI/frequency profiles, in particular, must be generated by a human operator. In general, the human operator modifies the target frequency profiles until the desired storage capacity is reached.

The format generator 68 includes a format efficiency process that uses the format optimizer 66 target high/low variable BPI format designs as well as the optimal predetermined number of high/low performing head allocations to modify and generate the appropriate configuration pages as part of the file system. For each zone, the format generator 68 selects the nearest frequency to the target frequency for that zone, given the component limitations mentioned above. The nearest frequency provides the target formats.

The optimal predetermined number of high/low performing head allocations comprises the number of heads allocated to each of the multiple frequencies in each zone. The format optimizer 66 determines the head allocation, which is input to the format generator 68. The capacity of a zone

depends on the target frequencies and the number of heads allocated to each frequency.

The format optimizer 66 uses the nominal average BPI or frequency (nominal BPI format target designs) (e.g., one read/write frequency) in each zone from the format generator 68 to estimate the yield before applying the variable BPI designs. For a design with multiple frequencies per zone, this is the weighted average by the number of allocated heads of the multiple frequencies. The nominal format is created by a human operator working with the format generator 68 in an interactive manner.

The format generator 68 calculates the number of tracks per zone, number of blocks per track, radius at each zone as well as block and track format efficiency to calculate the zone data storage. The format optimizer 66 estimates the zone data storage using the tracks per zone, radii, and format efficiency. Thus, the format optimizer 66 and the format generator 68 interact as shown in FIG. 2A. For example, in a disk drive with four heads, and two data density format frequency profiles (high and low frequency profiles) with three zones across the disk surface, after the measurement and optimization processes, the format generator 68 is provided with two optimal frequency profiles and the optimal allocation of the heads. The format generator 68 then calculates the storage capacity, and if the disk drive meets the minimum required storage capacity, the format generator 68 generates the configuration pages for the disk drive firmware. The configuration pages are used by the disk drive firmware to command the head to write at an assigned frequency to a zone. If the calculated storage capacity does not meet the minimum required storage capacity, the format optimization is performed again with new format efficiency values and the process is repeated.

Head Assignments

Allocating the number of heads to the predetermined multiple frequencies in a zone, and assigning a particular head in a particular disk drive to a particular frequency, are distinct. The allocation is performed by the format optimizer 66 and applies to the disk drives of a particular design. The head assignments are then performed during manufacturing as part of a test process undergone by each disk drive to be produced.

Once the configuration pages are generated and converted to binary files as part of the file system, they can be loaded into a reserved image of the disk drive for use after power cycling. Then, for every disk drive, the assignments are performed per head and per zone to assign a predetermined number of heads to high BPI formats and the remaining heads to low BPI formats in a two frequency design, to satisfy the allocation of heads to the formats by the format optimizer 66.

The head assignments for the two frequency format where high and low frequencies are used includes the steps of:

(1) Load default parameters from the configuration pages, and calibrate selected parameters on a per head, per zone basis (e.g., load high BPI format profile for all the zones across the stroke);

(2) Take measurements from the heads at the disk surfaces at preselected zones with respect to a target performance metric;

(3) For each head in every measured zone, sort/rank the heads by the target performance metric from best to worst, select a prespecified (by the allocation process in the format optimizer 66) number of heads with the best performance, and assign those heads to the high frequency for a particular zone;

(4) Optionally interpolate between the measurements obtained from the preselected number of zones to find the results for the other zones, and do the same for the interpolated zones. The interpolation reduces the test time. Head performances are measured, sorted and assigned to a frequency for a

subset of the total number of zones. For the remaining zones, the heads are assigned by interpolating the head assignments from the measurements;

(5) For every zone, save the worst prespecified number of weak heads with respect to the target performance metric; and

(6) For every zone, load and calibrate the weak heads with the low BPI format.

The above process can improve storage capacity, improve yield and tradeoff between storage capacity and yield. In a test, the heads can pass or fail with respect to a target performance metric to determine if the test target limits are met.

The disk drive firmware is extended to load more than one format profile. A head can be assigned a different read/write frequency per zone across a disk surface, and radially similarly situated zones on different disk surfaces can have different read/write frequencies assigned to the corresponding heads whereby one head is assigned a different frequency/format profile than another head.

The head assignments apply to a format design with two recording frequencies per zone, but can be easily extended to more than two frequencies per zone and can be iterated to assign heads to more than two frequencies per zone. For example, in a design with H heads and F frequencies per zone, steps 1 and 2 are completed for the high frequency. The first selection of heads in step 3 assigns the highest h1 heads, where h1 is the prespecified number of heads allocated to the highest frequency for that zone. The remaining (H - h1) heads are then loaded and calibrated with the second highest frequency (step 1 again), measurements are taken (step 2 again), the heads are ordered relative to the metric and the best h2 heads are assigned to the second highest frequency (step 3 again). Here h2 is the prespecified number of heads allocated to the second highest frequency in the zone. Steps 1-3 are then iterated for the (H - h1 - h2) heads, followed by the (H - h1 - h2 - h3) heads, and so on, until hF heads

remain to be assigned to the lowest frequency. The set of {h1 . . . hF} heads receive the head allocation made by the format optimizer 66.

Table 1 illustrates the vertical zoning head assignments on a disk drive with six heads and five zones across the stroke on each disk surface. Each head is assigned to either a high or low data density format based on record/playback performance of that head, and the number of heads assigned to high data density and the number of heads assigned to low data density is according to the head allocation determined by the format optimizer 66.

HEAD	ZONE 1	ZONE 2	ZONE 3	ZONE 4	ZONE 5
0	Low	High	Low	High	Low
1	High	Low	High	High	Low
2	High	Low	High	Low	High
3	High	High	Low	High	High
4	Low	High	High	High	High
5	High	High	High	Low	High

Table 1 – Example format assignment of a disk drive after test using vertical zoning with variable BPI across zones.

FIG. 3 shows a flowchart of vertical zoning data collection that includes the steps of:

(1) Select a number of disk drives for data measurement/collection (step 300);

(2) Create a nominal linear bit density profile KFCI (nominal KFCI): $\overline{\mathrm{kFCI}}\,(R)$, where R is the disk radius (step 302);

(3) Create more linear bit density profiles by multiplying the nominal KFCI by the scaling factor x_i (step 304):

$$(1 \pm x_i) * \overline{\mathrm{kFCI}}(R)$$

(4) Create a binary file system for every generated profile (step 306):

$$i \in \{1, \cdots, N\}$$

where N is the total number of frequency format profiles, for example, for N = 2, having X_1, and X_2, if $X_1 = 0.05$ and $X_2 = 0.1$, then including the nominal frequency format there are five different frequency profiles in step 304 as follows: (a) nominal KFCI, (b) 1.05 x nominal KFCI, (c) 0.95 x nominal KFCI, (d) 1.1 x nominal KFCI, and (e) 0.90 x nominal KFCI;

(5) Select the first head by setting *i* to 1(step 308);

(6) Load the file system *i* into the reserved image of the disk drives (step 310);

(7) Take the head performance measurements (step 312);

(8) Unload and save the results in the data base (step 314);

(9) Increment *i* by one (step 316);

(10) Determine if $i = N$ (step 318);

(11) If not, go to step 310; and

(12) Otherwise, stop (step 320).

The above process collects performance data for all the heads at all the zones.

FIG. 4 shows a flowchart of vertical zoning post-measurement and per zone BPI distribution extraction that includes the steps of:

(1) Organize the head performance data for every head $i \in \{1, \cdots, M_1\}$ and every zone $j \in \{1, \cdots, M_2\}$ as a function of the linear bit density samples, where M_1 is the total number of heads in the disk drives selected for measurement and M_2 is the total number of zones, to generate head performance histograms (step 400);

(2) Choose a target performance metric (step 402);

(3) Set $j = 1$ and $i = 1$ (step 404);

(4) Interpolate/extrapolate BPI at the target performance metric for head i at zone j (step 406);

(5) Select the next head by incrementing i by one (step 408);

(6) Determine if all the heads have been processed by determining if $i = M_1$ (step 410);

(7) If not, go to step 406 to process the next head, otherwise generate a frequency capability histogram at zone j for all the heads (step 412);

(8) Determine if all the zones have been processed by determining if $j = M_2$ (step 414);

(9) If not, move to the next zone and start with the first head again, set $j = j+1$ and $i = 1$ (step 416) and go to step 406; and

(10) Otherwise, stop (step 418).

The above process generates 1D frequency capability histograms at a target performance metric for every zone by considering all the heads from the sample disk drives selected for measurement. Using the 1D frequency capability histograms at a target performance metric, probability theory known to those skilled in the art can be adopted to estimate the 1D frequency capability distributions. Further, the above process is extended by using 2D and 3D interpolation/extrapolation routines to extract and estimate the 2D and 3D joint frequency capability histograms and their associated distributions.

FIG. 5 shows a flowchart of head assignments for N heads with a two frequency format (high/low data density) that includes the steps of:

(1) Assign all the heads in a disk drive to the high data density format (step 500);

(2) Calibrate all the heads at the high data density format for selected zones (step 502);

(3) Measure the head performance metric at the selected zones for all the heads (step 504);

(4) For each selected zone, rank the heads by the head performance metric (step 506);

(5) For each selected zone, assign the highest K heads to the high data density format, and assign the other N-K heads to the low data density format (step 508);

(6) Optionally interpolate the head assignments for the remaining zones (step 510); and

(7) Complete the calibration of all the heads and all the zones at the assigned formats (step 512).

The above process completes the assignment of each head in each disk drive to a predetermined frequency.

The format generator 68 passes the track formats to the format optimizer 66 to have a more accurate way of calculating the storage capacity (nominal format). Such information and constraints are provided to the format optimizer 66 to solve the joint optimization problems. The format optimizer 66 performs a coarse calculation of the storage capacity, whereas the format generator 68 performs an exact calculation of the storage capacity. The format generator 68 provides format information (such as number of tracks per zone, and the zone format) to the format optimizer 66, and calculates the exact storage capacity. Such information is passed once from the format generator 68 to the format optimizer 66 for a head design with a given number of heads. The format generator 68 initially provides nominal information to the format optimizer 66, and the format optimizer 66 performs its calculation of target densities (zone frequencies and number of heads allocated to each frequency) and provides that information to the format generator 68. The format generator 68 then determines if required storage capacity has been reached. Adjusting the target densities to meet storage capacity and/or yield requirements includes adjusting the selected zone density or zone frequencies.

FIG. 6 shows a flowchart of format generation and optimization in which an iterative process for a minimum storage capacity (C) and a user specified storage overcapacity (Δ) includes the steps of:

(1) Determine the disk geometry, track density and servo spoke details, and provide the inner diameter (ID) and outer diameter (OD) radii, the track density (TPI), the number of servo spokes, and the servo spoke length (step 600);

(2) The format generator 68 generates the initial format at the storage capacity using the ID and OD radii, the TPI, the number of servo spokes and the servo spoke length, and provides the radius of each zone per disk surface, the number of tracks per zone, the number of blocks per track, and the format efficiency by zone (step 602);

(3) The format optimizer 66 generates optimal target densities at all the zones using the radius of each zone per disk surface, the number of tracks per zone, the number of blocks per track, and the format efficiency by zone and provides the high and low BPI targets by zone and the number of high and low BPI head allocations by zone (step 604);

(4) The format generator 68 generates new formats with a storage capacity (the number of logical blocks per disk drive) (step 606);

(5) Determine if the storage capacity > C and the storage capacity < (C + Δ) (step 608);

(6) If not, adjust the target densities (step 610) and go to step 606; and

(7) Otherwise, stop (step 612).

For example, the disk surface capacity is TPI x BPI x (1 + ECC) / FE, where TPI is the track density, BPI is the linear bit density, ECC is the fractional level of error correcting code which is typically about 0.1, and FE is the format efficiency which is typically about 0.57.

The above process completes the format generation.

As another example, a set of thirty-two mature disk drives are selected and each disk drive includes twelve heads. The 1D, 2D and 3D BPI distributions are extracted at an on-track symbol error rate from the outer, middle and inner zones. Next, the BPI distributions are fed to the format optimizer 66, and high or low frequency per zone format designs are obtained at the three zones. This is performed once by individual optimization based on 1D BPI distributions at each of the three zones, and once by joint optimization based on the measurements obtained from the three zones and their extracted 1D, 2D and 3D BPI distributions. The head format allocation search is performed by simulation, and for each zone the one-format designs before the application of vertical zoning are a special case of the two-format variable BPI designs by forcing the high and low formats to be equal to the nominal BPI format at that zone. Furthermore, the pass/fail of the disk drives is based on each head at every zone passing a target on-track symbol error rate as well as off-track squeeze and unsqueeze offset margins. Then, the yield is calculated by simulation by interpolation/extrapolation of the measurement data before and after the application of vertical zoning (VZ). Table 2 summarizes the results:

	Using Joint Optimization	Using Individual Optimization
Drive Yield (Yd)	93.75	90.625
Drives failed after VZ	4 & 29	4, 6 & 29
Drives recovered	2, 3, 13, 19, 21 & 25	2, 3, 13, 19, 21 & 25
Passed drives failed after VZ	None	6
Drives failed before VZ	2, 3, 4, 12, 19, 21 & 29 i.e., drive yield before VZ Yd=75%	2, 3, 4, 12, 19, 21 & 29 i.e., drive yield after VZ Yd = 75%

Table 2

Although a manufacturing test case for a two format design is illustrated, the search algorithm can be easily generalized to a higher number of formats. The design of two formats based on 1D, 2D and 3D BPI distributions can easily be generalized to higher order or dimensions by considering more than three zones. The format design can be generalized from two to a higher number of formats. The measurement procedure can be generalized to consider more zones as well as off-track measurements such as 747 curves or quality metrics

versus error rate measurements to perform a correlation study for the best metric with the least test time.

Further, the per zone variable BPI design can be easily extended to a variable BPI/TPI design. The measurement process is extended to include 747 measurements of all the heads from a preselected number of disk drives. To speed up the measurements, instead of 747 measurements, off-track and adjacency margin squeeze measurements of the heads can be performed. Once the 747 data of the heads at a preselected number of zones is determined, for every zone, joint BPI/TPI distributions can be extracted at the target(s) by post-measurement data processing. The choice of a target is an integral part of the performance gain, such as yield, due to the per zone variable BPI/TPI designs. Some example targets are off-track symbol error rate, position error signal variance, and a combination of both. After the joint BPI/TPI distributions are extracted and available for the zones, a per zone variable BPI/TPI design can be obtained by solving two constrained (joint) optimization problems: one that maximizes the yield while keeping the same disk drive areal density, and another that maximizes the disk drive areal density while keeping the same yield. Once the per zone variable BPI/TPI designs are obtained, a preselected number of heads are allocated to high and low density BPI and TPI formats, for example, for a two variable BPI/TPI per zone design performed as part of the test process.

The present invention improves storage capacity (and consequently areal density at a fixed target BPI) and yield and reduces the target TPI by increasing the average BPI across the stroke per head (depending on the number of formats considered) to meet a desire storage capacity. In particular, the BPI at the outer diameter may be limited by the maximum deliverable data rate of the ASIC components. For example, if the controller 57 has a maximum deliverable data rate of 650 MHz, the preamplifier 21 has a maximum deliverable data rate of 700 MHz and the read/write channel 51 has a maximum deliverable data rate of 750 MHz, then the BPI at the outer diameter is limited by the controller 57 at a

maximum deliverable data rate of 650 MHz. Thus, a conventional one format BPI profile across the stroke does not achieve the desired storage capacity and yield, whereas the present invention per zone variable BPI target formats meet the desired storage capacity at a fixed target TPI while improving the yield.

The data measurer 62 can be implemented by a general purpose computer 61 and the drive electronics of the disk drive 100. The general purpose computer 61 can be a high end PC, a PC server or a workstation and include programmable simulation software. The drive electronics can include the logic circuit 49 and the controller 57. The logic circuit 49 performs the measurements and the controller 57 directs the logic circuit 49 and transfers the data to the general purpose computer 61. The head assignments can be implemented by the controller 57 with the data collection sub-task performed by the logic circuit 49. The post-measurement data processor 64, the format optimizer 66 and the format generator 68 can be implemented by the general purpose computer 61.

The present invention has been described in considerable detail with reference to certain preferred versions thereof; however, other versions are possible. Therefore, the spirit and scope of the appended claims should not be limited to the description of the preferred versions contained herein.

Abstract

A method of defining a storage format in data storage devices, each data storage device having multiple storage media and corresponding heads, each head for recording on and playback of information from a corresponding storage media in at least one zone, and each zone including concentric tracks for recording on and playback of information. The method includes selecting a sample of the data storage devices, for each selected data storage device measuring a record/playback performance capability of each head at one or more read/write frequencies per zone, generating storage density distributions corresponding to the heads in the selected data storage devices based on the performance capability measurements, selecting a group of read/write frequencies for the data storage devices with two or more frequencies for each zone based on the storage density distributions, and assigning one of the read/write frequencies to each head based on the performance capability of that head.

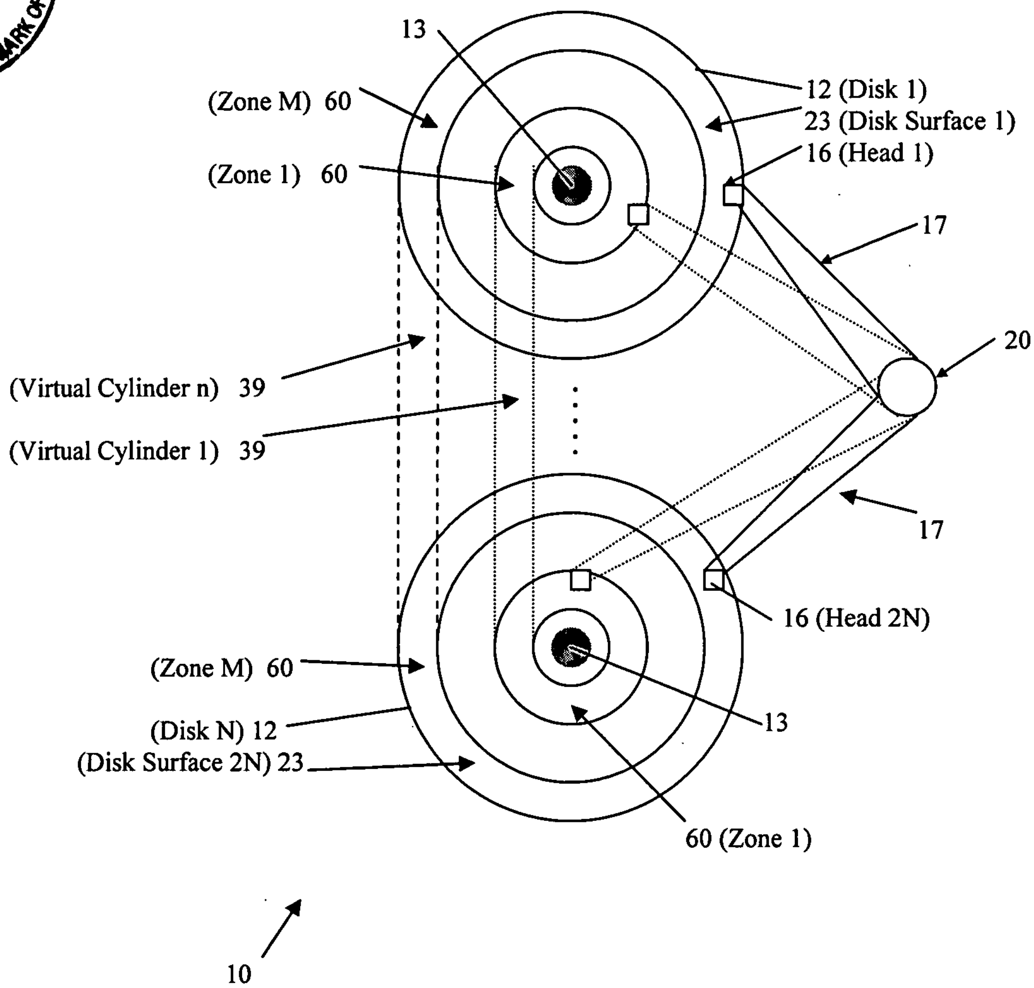
13
(Zone M) 60
(Zone 1) 60
12 (Disk 1)
23 (Disk Surface 1)
16 (Head 1)
17
20
(Virtual Cylinder n) 39
(Virtual Cylinder 1) 39
17
16 (Head 2N)
(Zone M) 60
13
(Disk N) 12
(Disk Surface 2N) 23
60 (Zone 1)
10

FIG. 1A

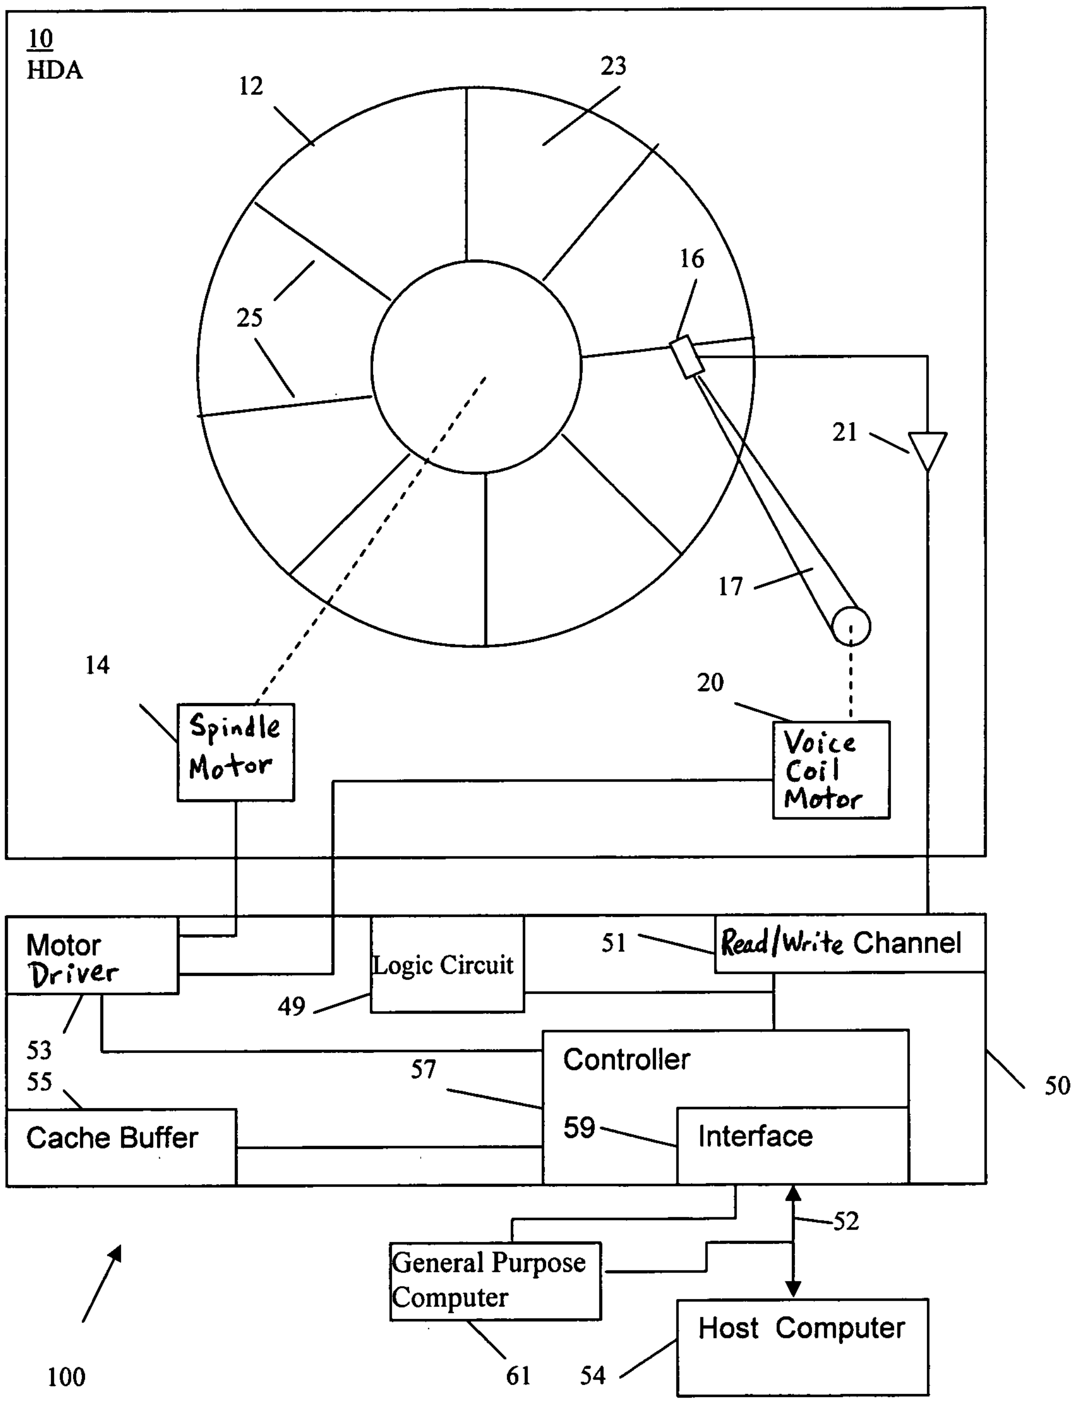
10
HDA
12
23
16
25
21
17
14
20
Spindle Motor
Voice Coil Motor
Motor Driver
Logic Circuit
51
Read/Write Channel
49
53
55
57
Controller
50
59
Interface
Cache Buffer
52
General Purpose Computer
Host Computer
100
61
54

FIG. 1B

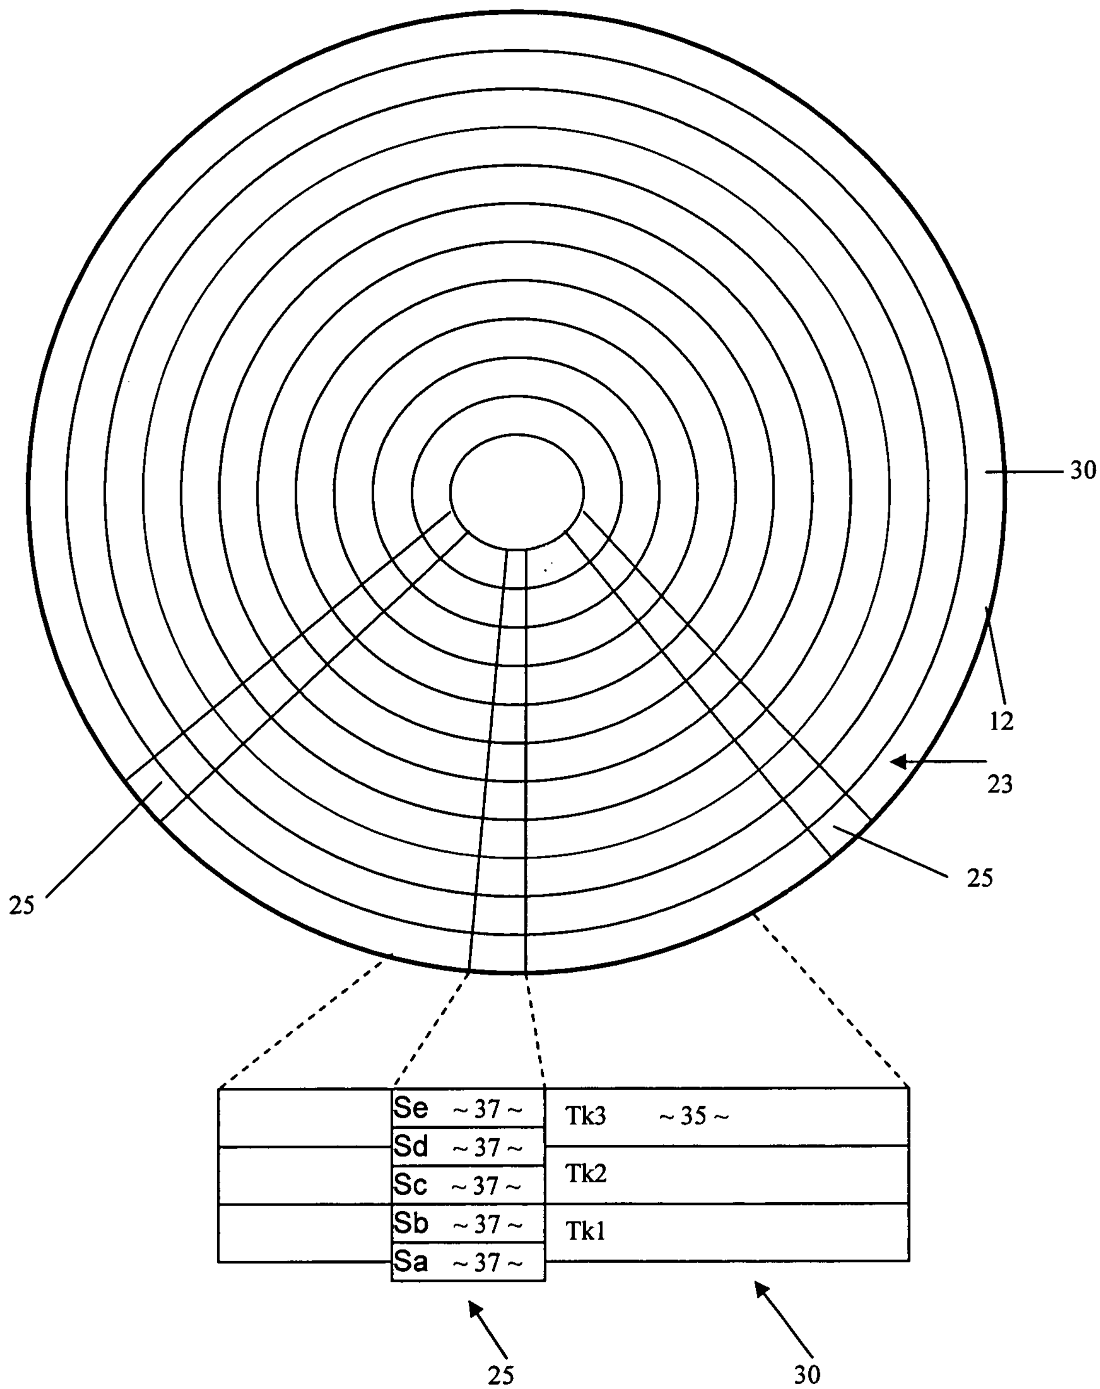

FIG. 1C

FIG. 1D

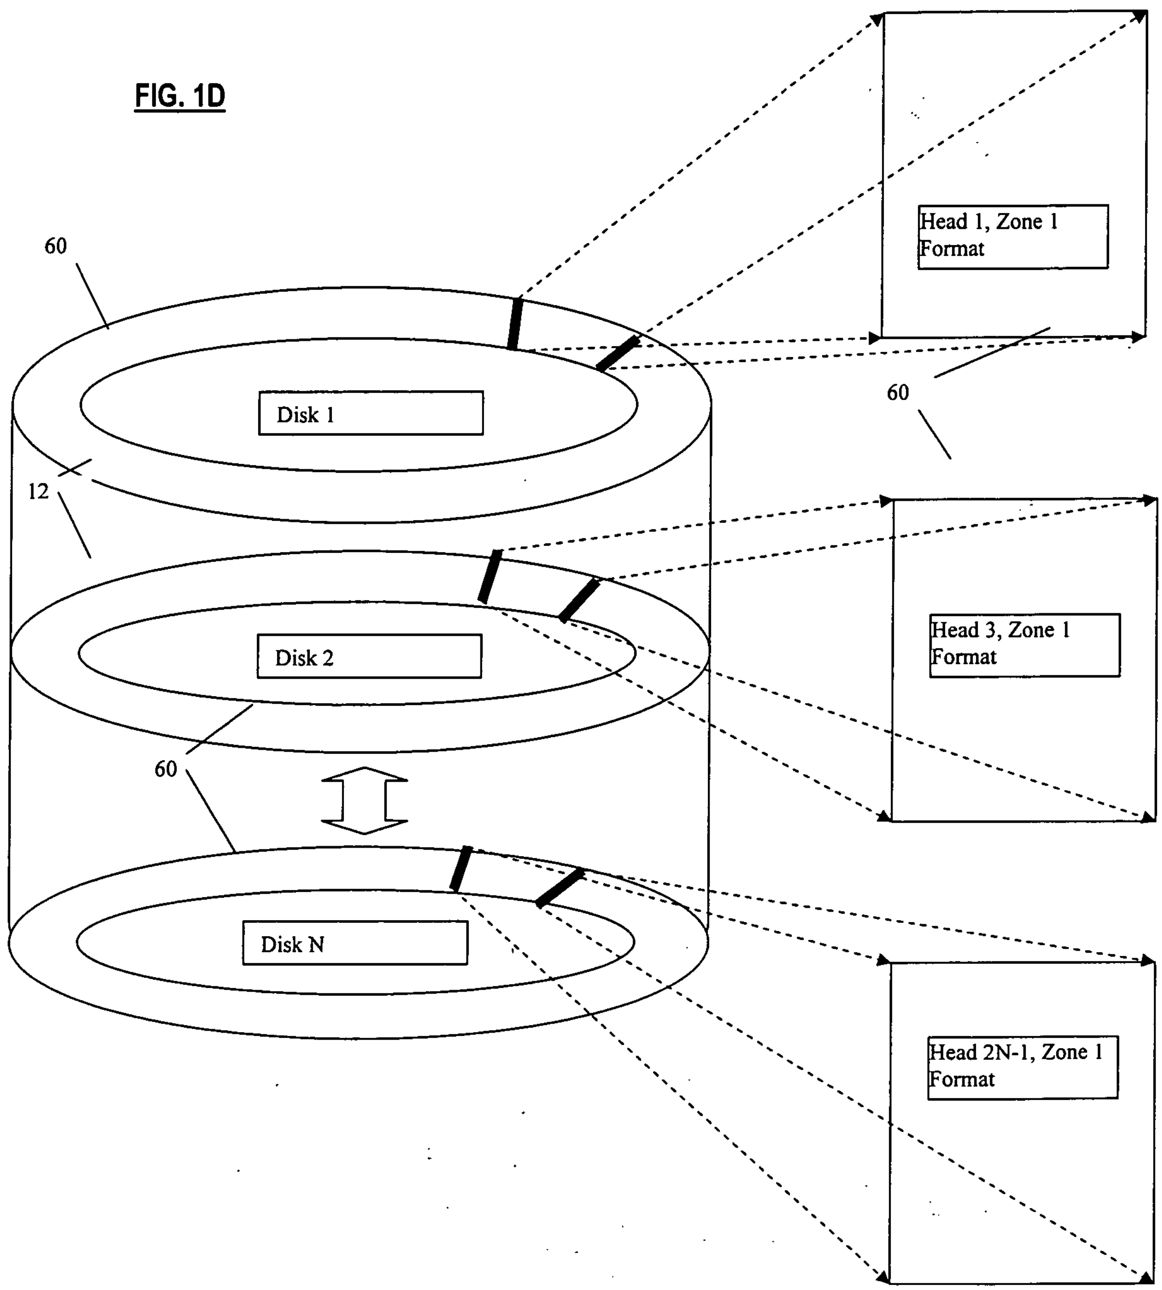
60
12
Disk 1
Disk 2
60
Disk N
Head 1, Zone 1
Format
60
Head 3, Zone 1
Format
Head 2N-1, Zone 1
Format

FIG. 1E

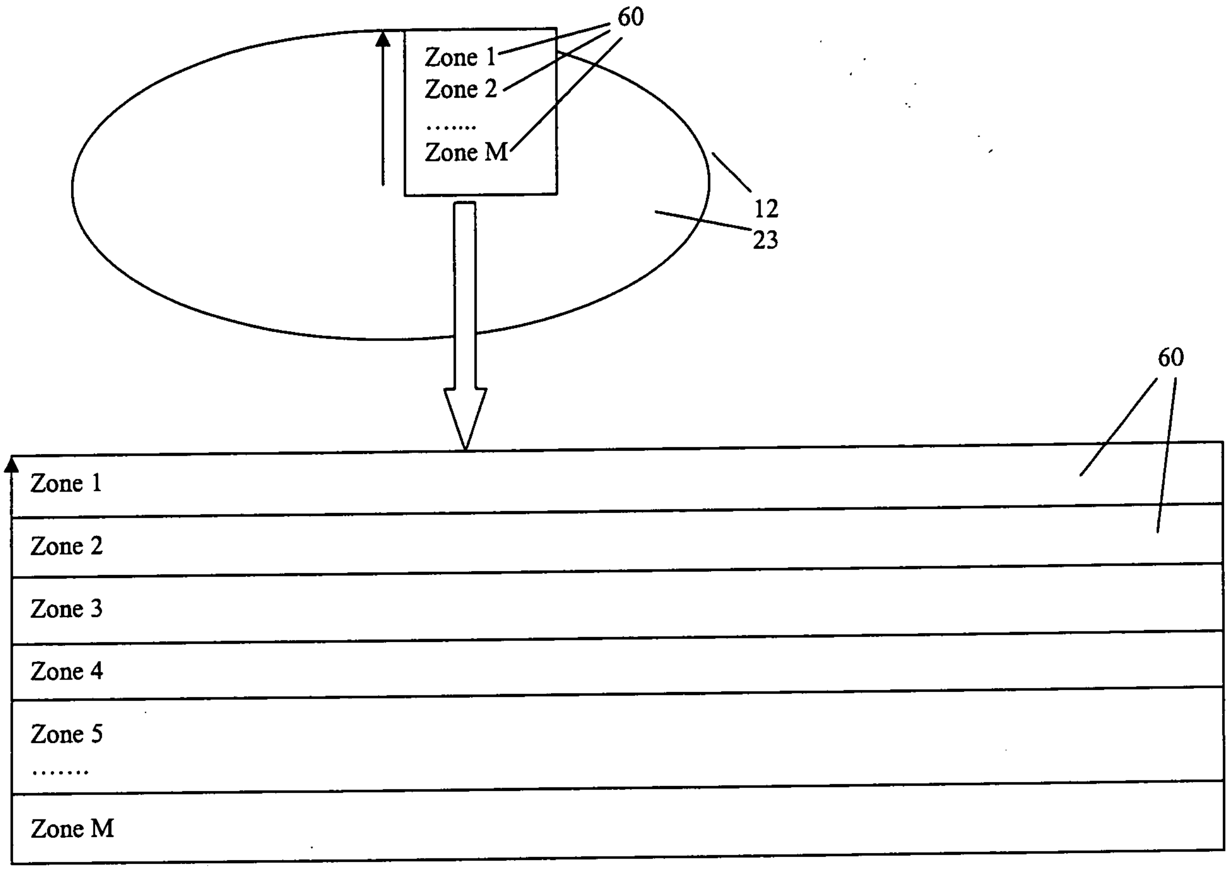
60
Zone 1
Zone 2
....
Zone M
12
23
60
Zone 1
Zone 2
Zone 3
Zone 4
Zone 5
....
Zone M

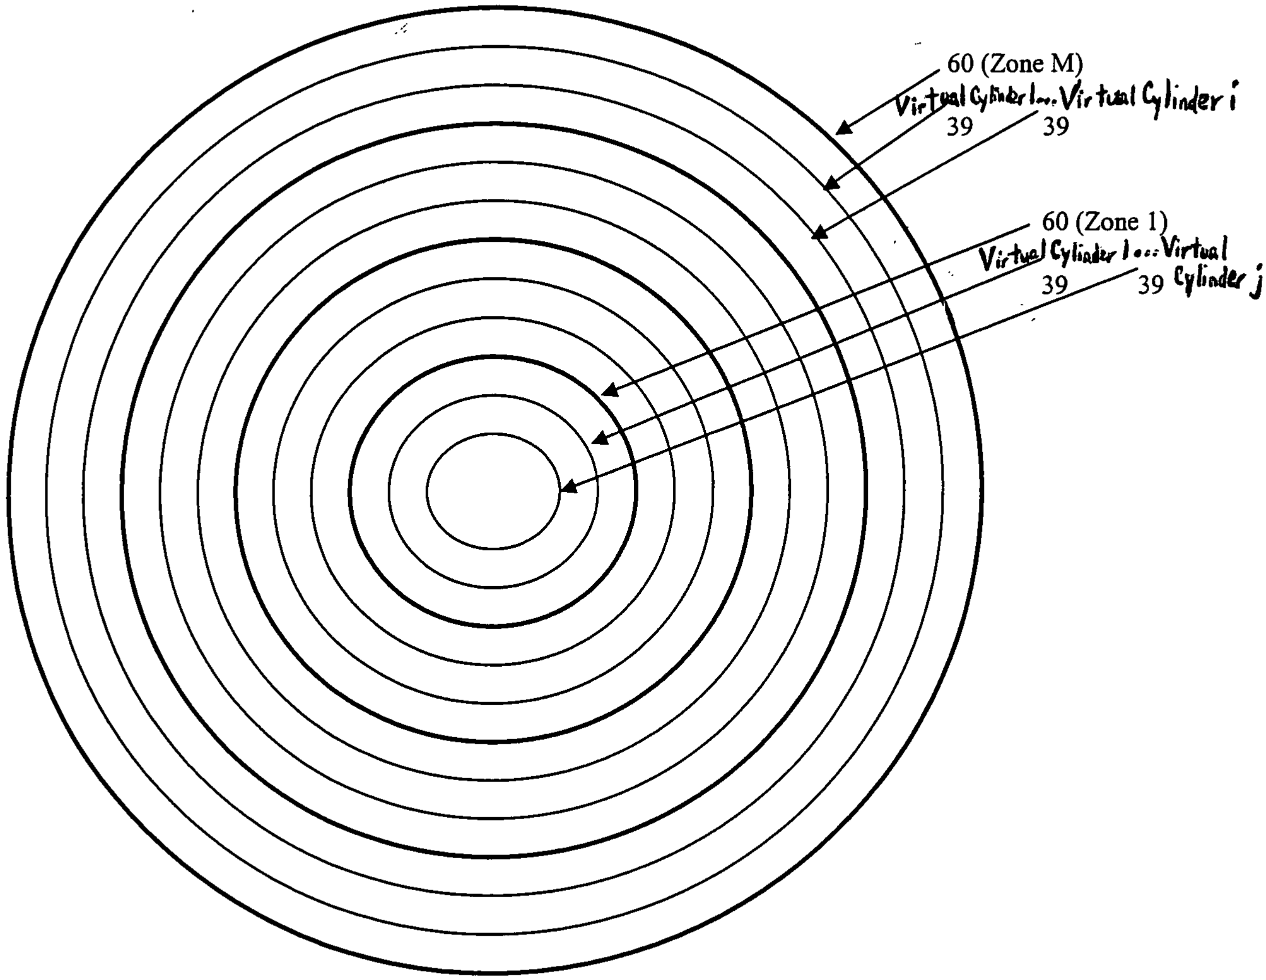
60 (Zone M)
Virtual Cylinder 1... Virtual Cylinder i
39
39
60 (Zone 1)
Virtual Cylinder 1... Virtual Cylinder j
39
39

FIG. 1F

FIG. 1G

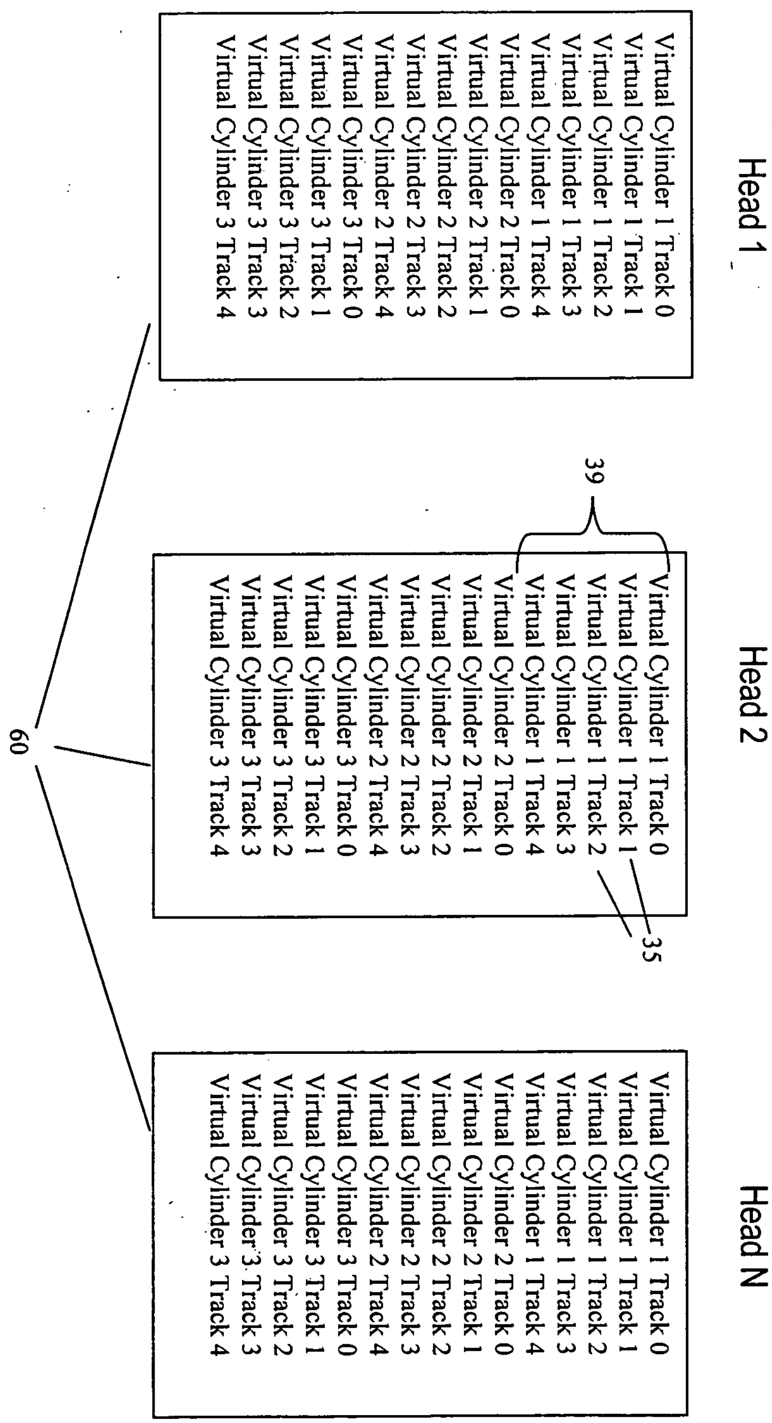
Head 1
Virtual Cylinder 1 Track 0
Virtual Cylinder 1 Track 1
Virtual Cylinder 1 Track 2
Virtual Cylinder 1 Track 3
Virtual Cylinder 1 Track 4
Virtual Cylinder 2 Track 0
Virtual Cylinder 2 Track 1
Virtual Cylinder 2 Track 2
Virtual Cylinder 2 Track 3
Virtual Cylinder 2 Track 4
Virtual Cylinder 3 Track 0
Virtual Cylinder 3 Track 1
Virtual Cylinder 3 Track 2
Virtual Cylinder 3 Track 3
Virtual Cylinder 3 Track 4
Head 2
39
Virtual Cylinder 1 Track 0
Virtual Cylinder 1 Track 1
Virtual Cylinder 1 Track 2
Virtual Cylinder 1 Track 3
Virtual Cylinder 1 Track 4
Virtual Cylinder 2 Track 0
Virtual Cylinder 2 Track 1
Virtual Cylinder 2 Track 2
Virtual Cylinder 2 Track 3
Virtual Cylinder 2 Track 4
Virtual Cylinder 3 Track 0
Virtual Cylinder 3 Track 1
Virtual Cylinder 3 Track 2
Virtual Cylinder 3 Track 3
Virtual Cylinder 3 Track 4
35
60
Head N
Virtual Cylinder 1 Track 0
Virtual Cylinder 1 Track 1
Virtual Cylinder 1 Track 2
Virtual Cylinder 1 Track 3
Virtual Cylinder 1 Track 4
Virtual Cylinder 2 Track 0
Virtual Cylinder 2 Track 1
Virtual Cylinder 2 Track 2
Virtual Cylinder 2 Track 3
Virtual Cylinder 2 Track 4
Virtual Cylinder 3 Track 0
Virtual Cylinder 3 Track 1
Virtual Cylinder 3 Track 2
Virtual Cylinder 3 Track 3
Virtual Cylinder 3 Track 4

FIG. 1H

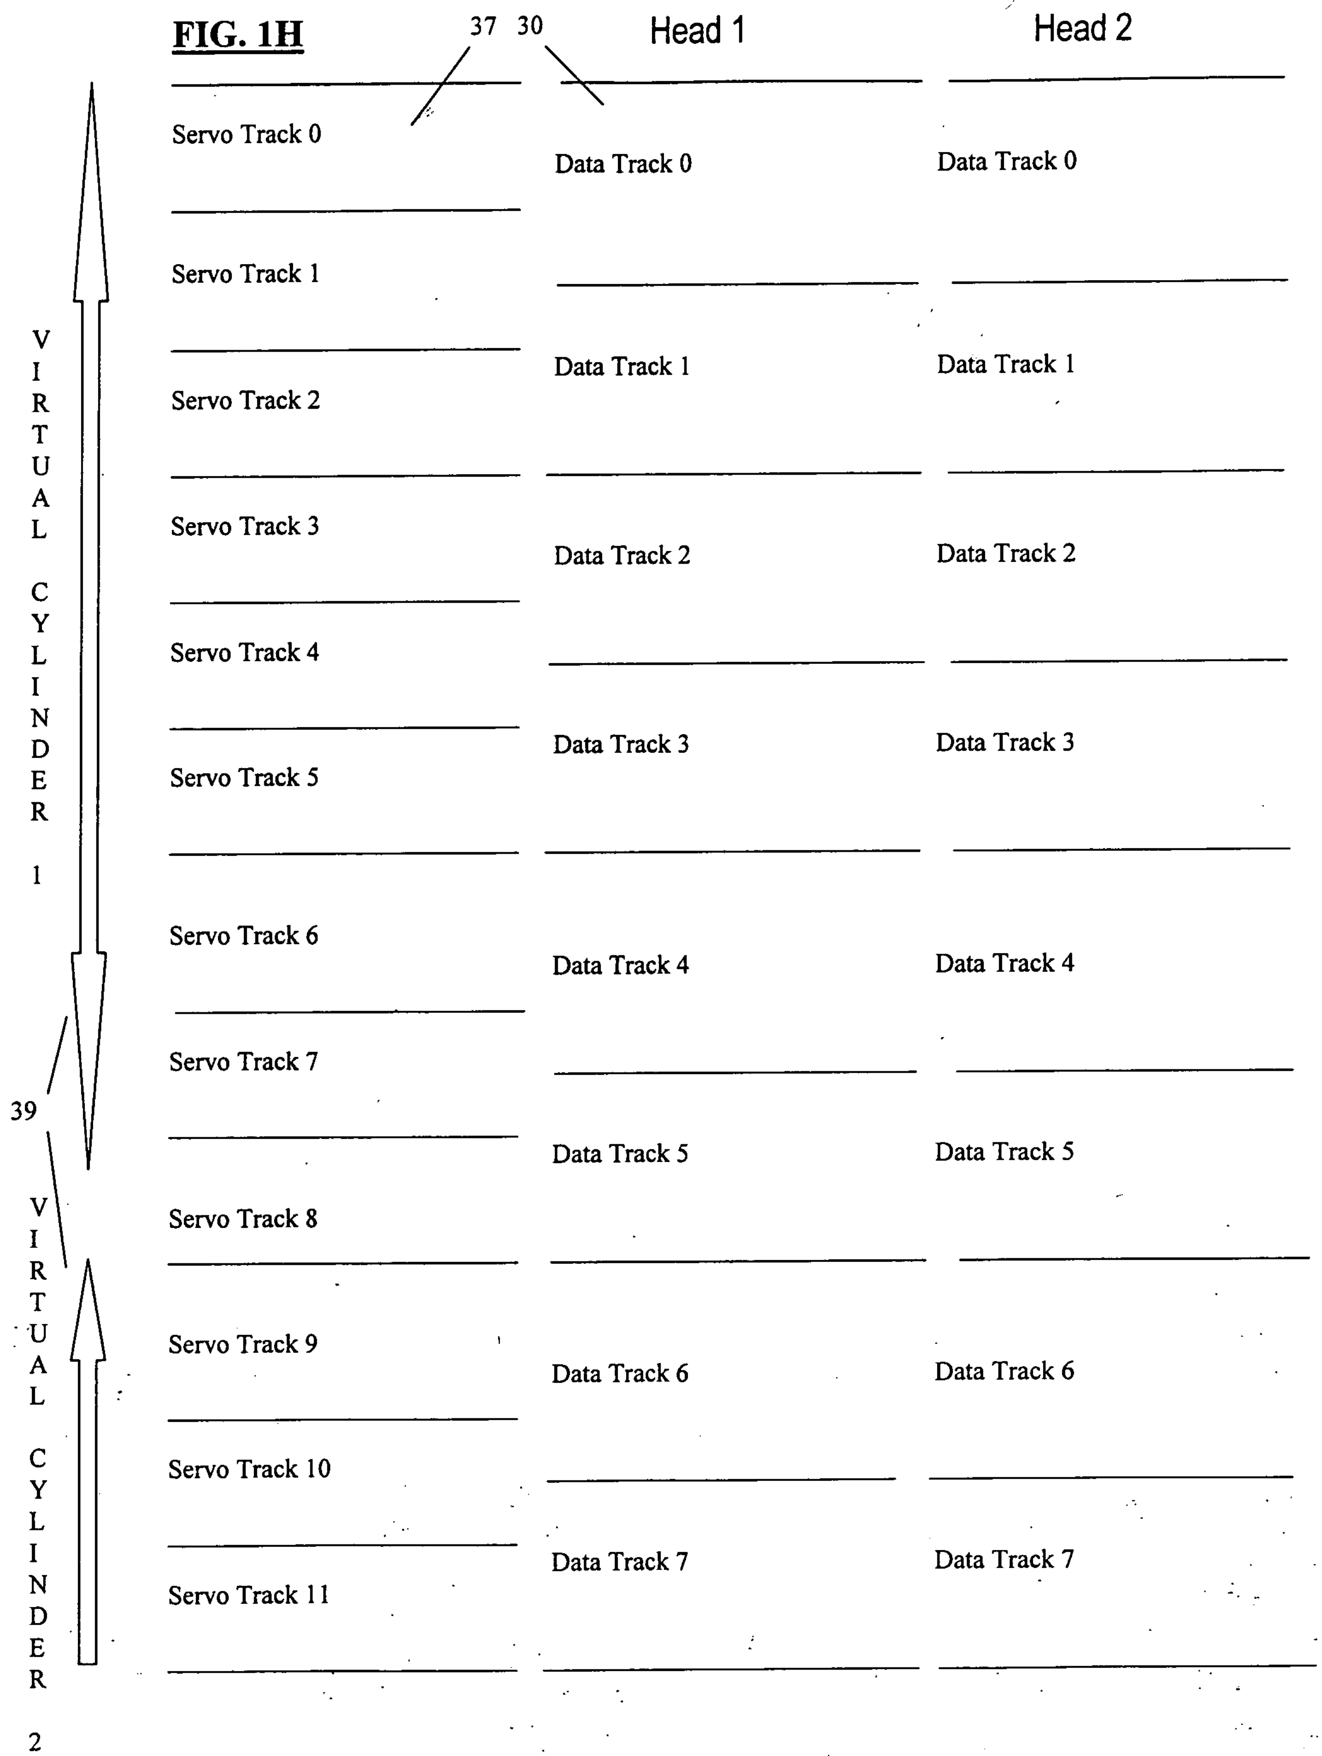
37
30
Head 1
Head 2
Servo Track 0
Servo Track 1
Servo Track 2
Servo Track 3
Servo Track 4
Servo Track 5
Servo Track 6
Servo Track 7
Servo Track 8
Servo Track 9
Servo Track 10
Servo Track 11
Data Track 0
Data Track 1
Data Track 2
Data Track 3
Data Track 4
Data Track 5
Data Track 6
Data Track 7
Data Track 0
Data Track 1
Data Track 2
Data Track 3
Data Track 4
Data Track 5
Data Track 6
Data Track 7
VIRTUAL CYLINDER 1
39
VIRTUAL CYLINDER 2

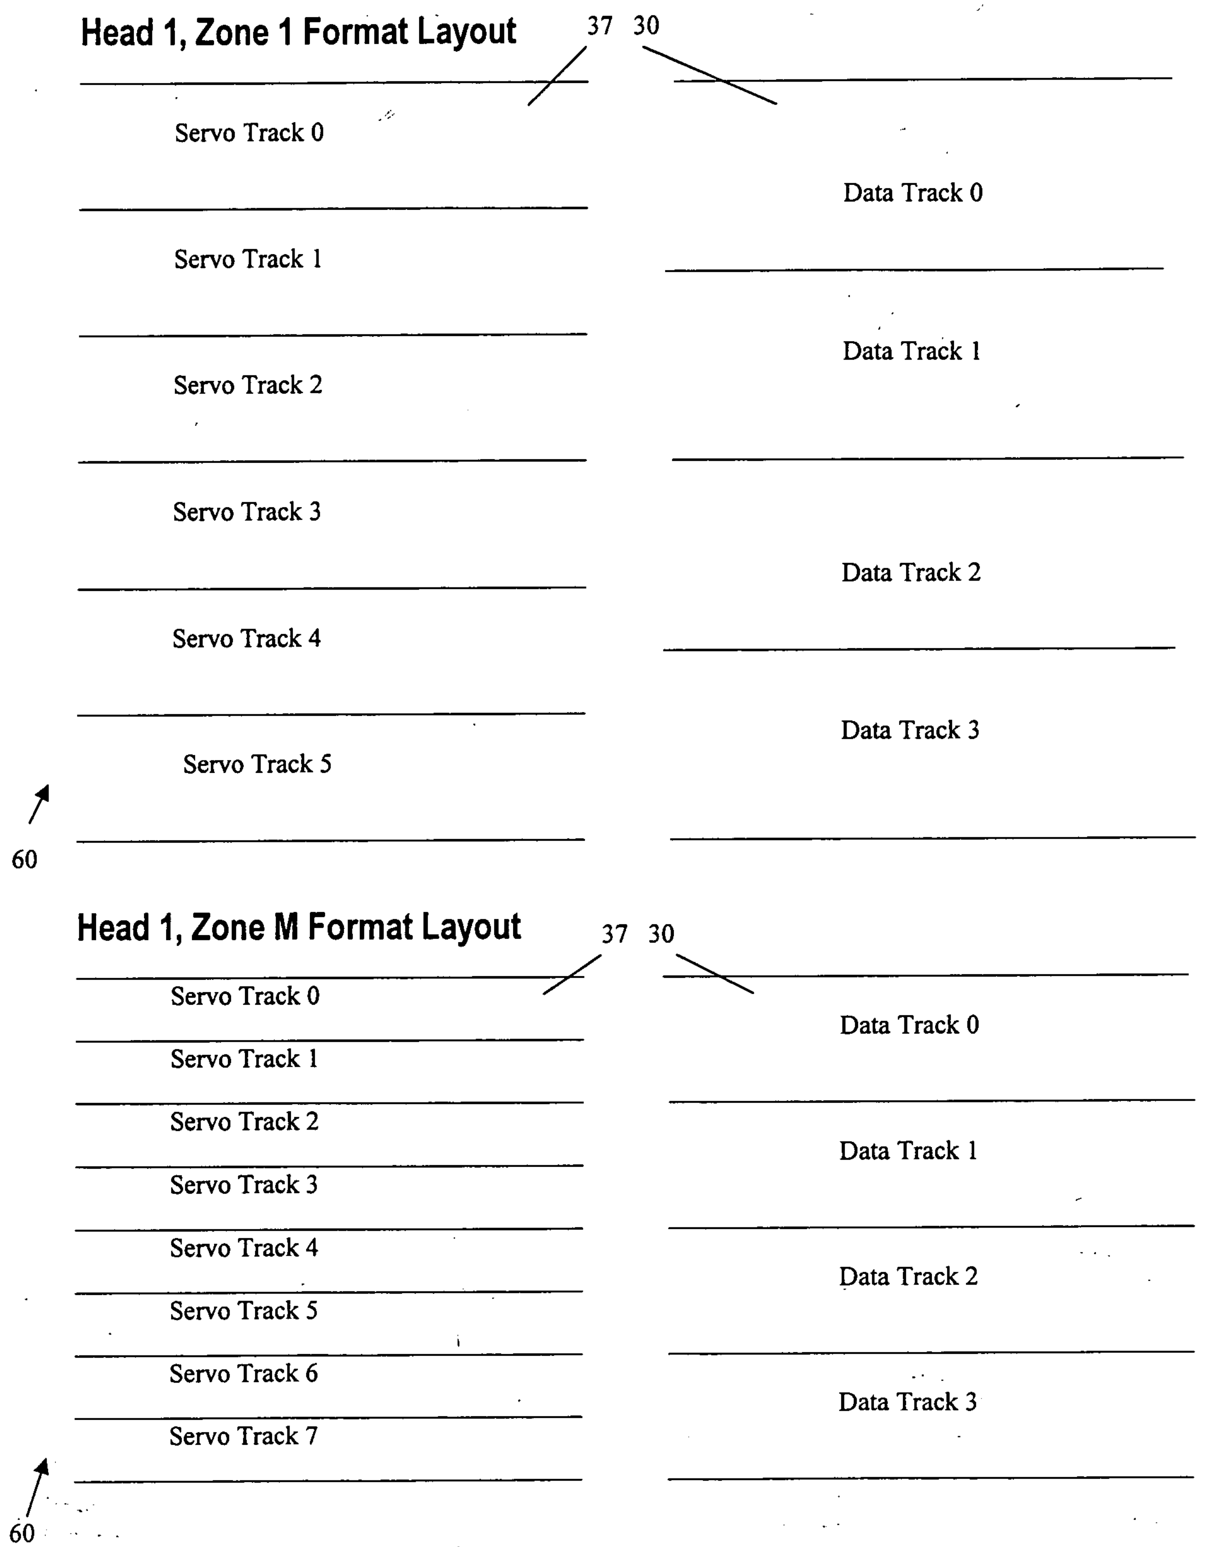
Head 1, Zone 1 Format Layout
37
30
Servo Track 0
Servo Track 1
Servo Track 2
Servo Track 3
Servo Track 4
Servo Track 5
Data Track 0
Data Track 1
Data Track 2
Data Track 3
60
Head 1, Zone M Format Layout
37
30
Servo Track 0
Servo Track 1
Servo Track 2
Servo Track 3
Servo Track 4
Servo Track 5
Servo Track 6
Servo Track 7
Data Track 0
Data Track 1
Data Track 2
Data Track 3
60

FIG. 1I

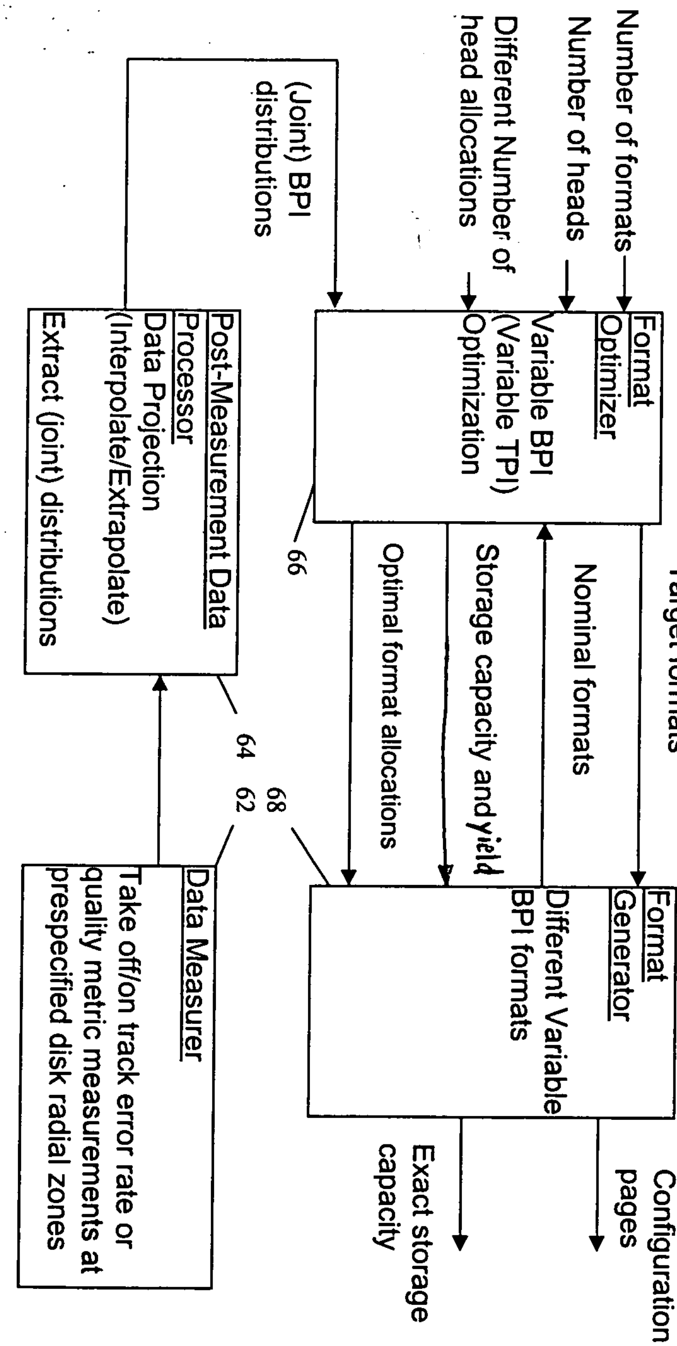
Number of formats
Number of heads
Different Number of head allocations
(Joint) BPI distributions
Format Optimizer
Variable BPI (Variable TPI) Optimization
66
Target formats
Nominal formats
Storage capacity and yield
Optimal format allocations
Format Generator
Different Variable BPI formats
Configuration pages
Exact storage capacity
68
62
64
Post-Measurement Data Processor
Data Projection (Interpolate/Extrapolate)
Extract (joint) distributions
Data Measurer
Take off/on track error rate or quality metric measurements at prespecified disk radial zones

FIG. 2A

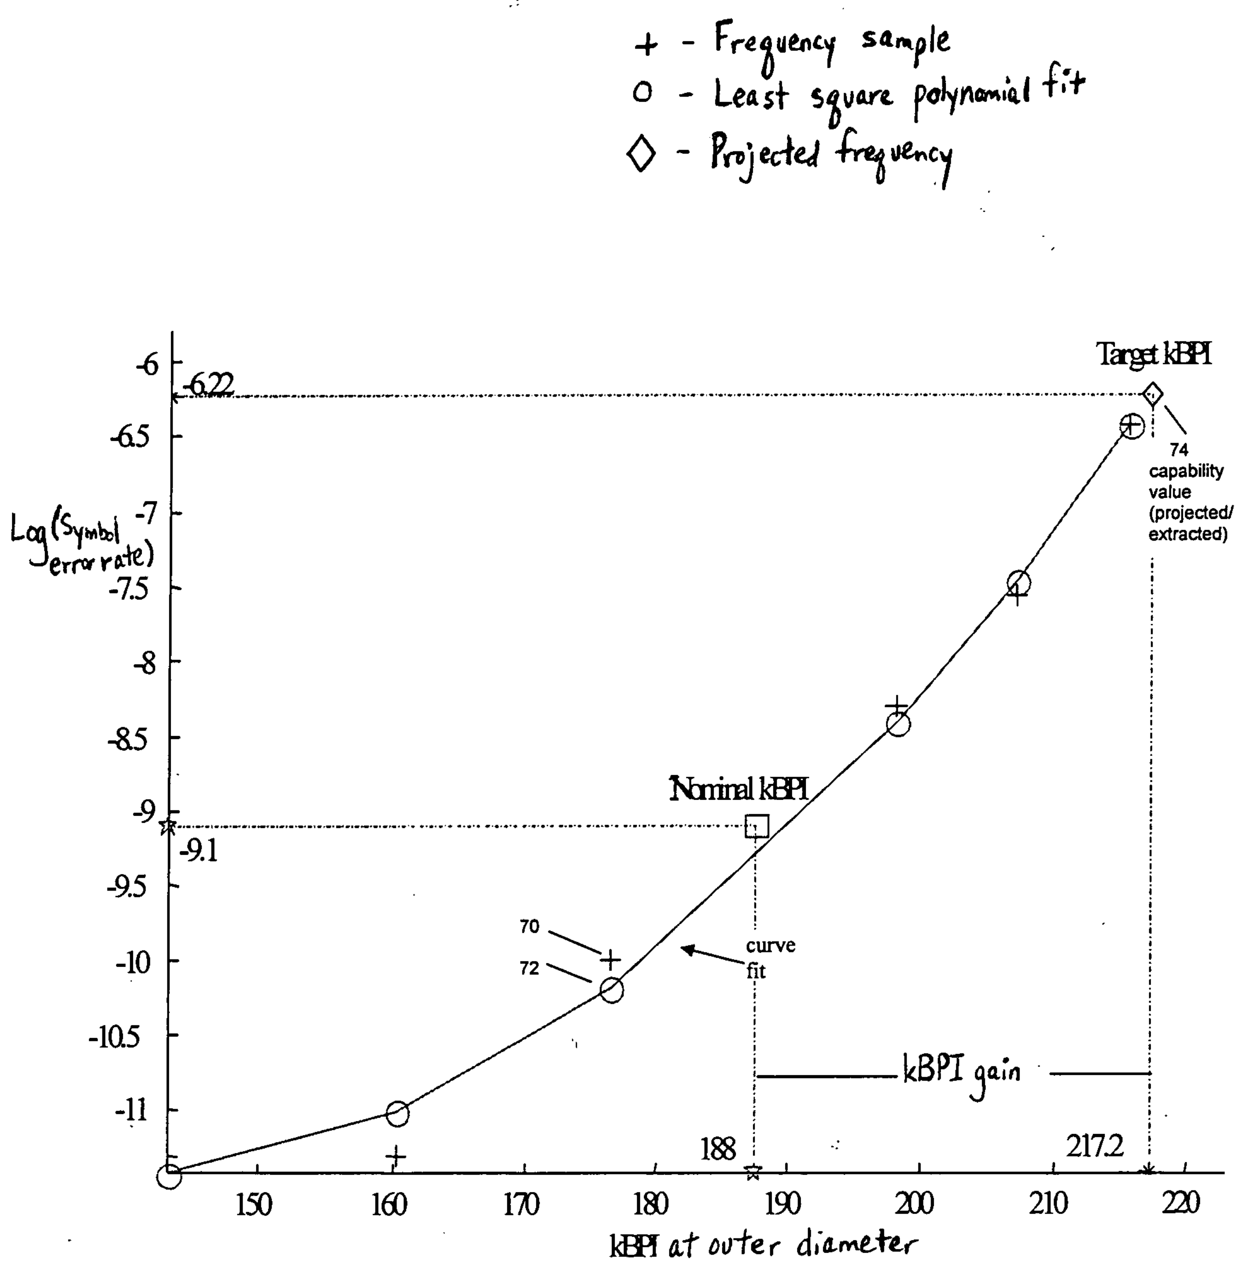
+ - Frequency sample
o - Least square polynomial fit
◇ - Projected frequency
Target kBPI
-6.22
74
capability value (projected/ extracted)
Log(Symbol error rate)
-6
-6.5
-7
-7.5
-8
-8.5
-9
-9.1
-9.5
-10
-10.5
-11
Nominal kBPI
70
72
curve fit
kBPI gain
188
217.2
150
160
170
180
190
200
210
220
kBPI at outer diameter

FIG. 2B

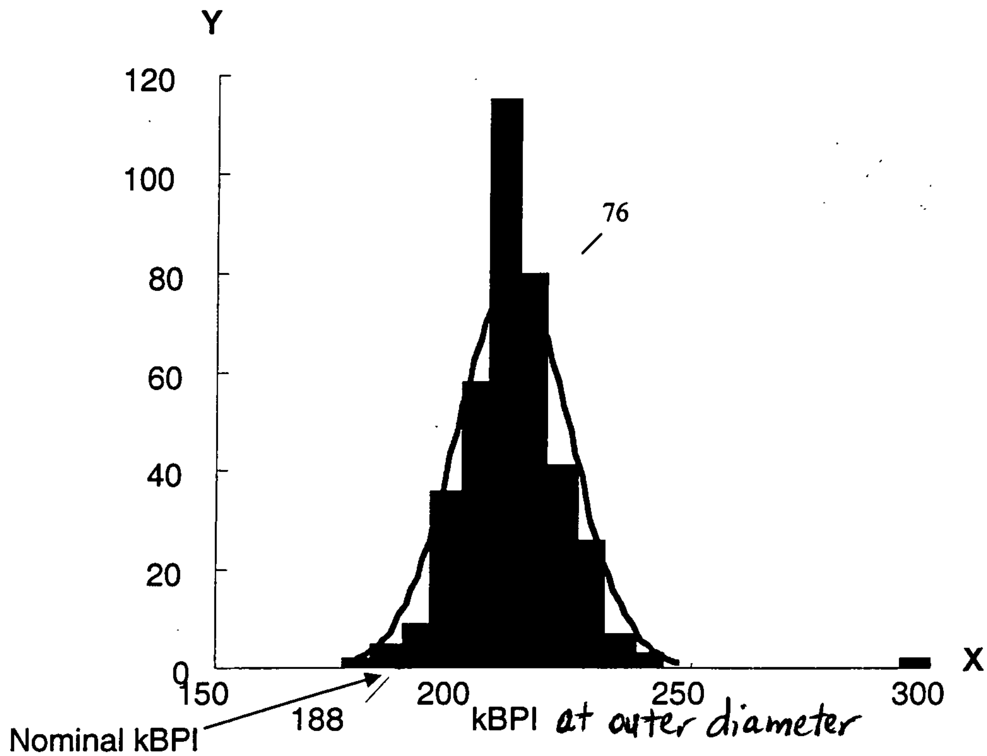
Y
120
100
80
60
40
20
0
76
150
188
200
250
300
X
kBPI at outer diameter
Nominal kBPI

FIG. 2C

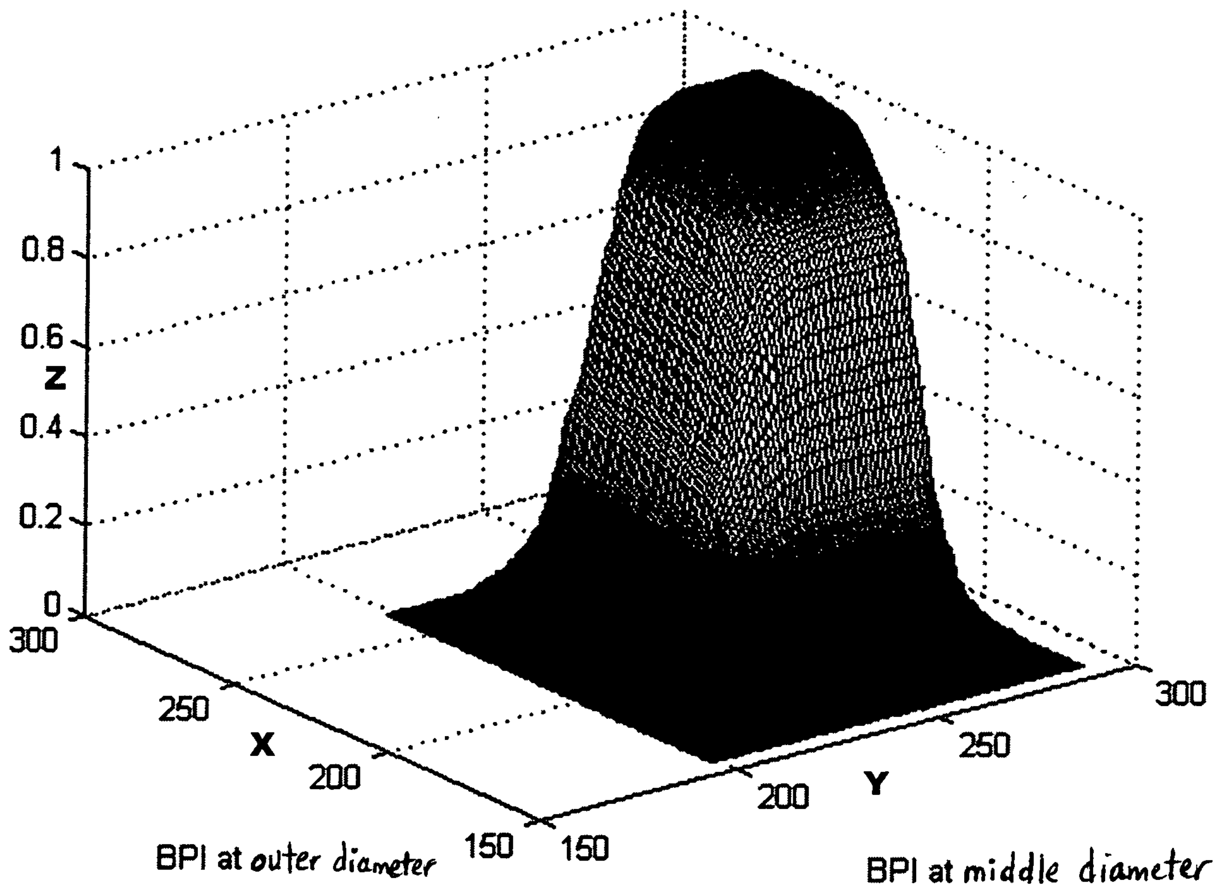
1
0.8
0.6
Z
0.4
0.2
0
300
250
X
200
150
BPI at outer diameter
150
200
Y
250
300
BPI at middle diameter

FIG. 2D

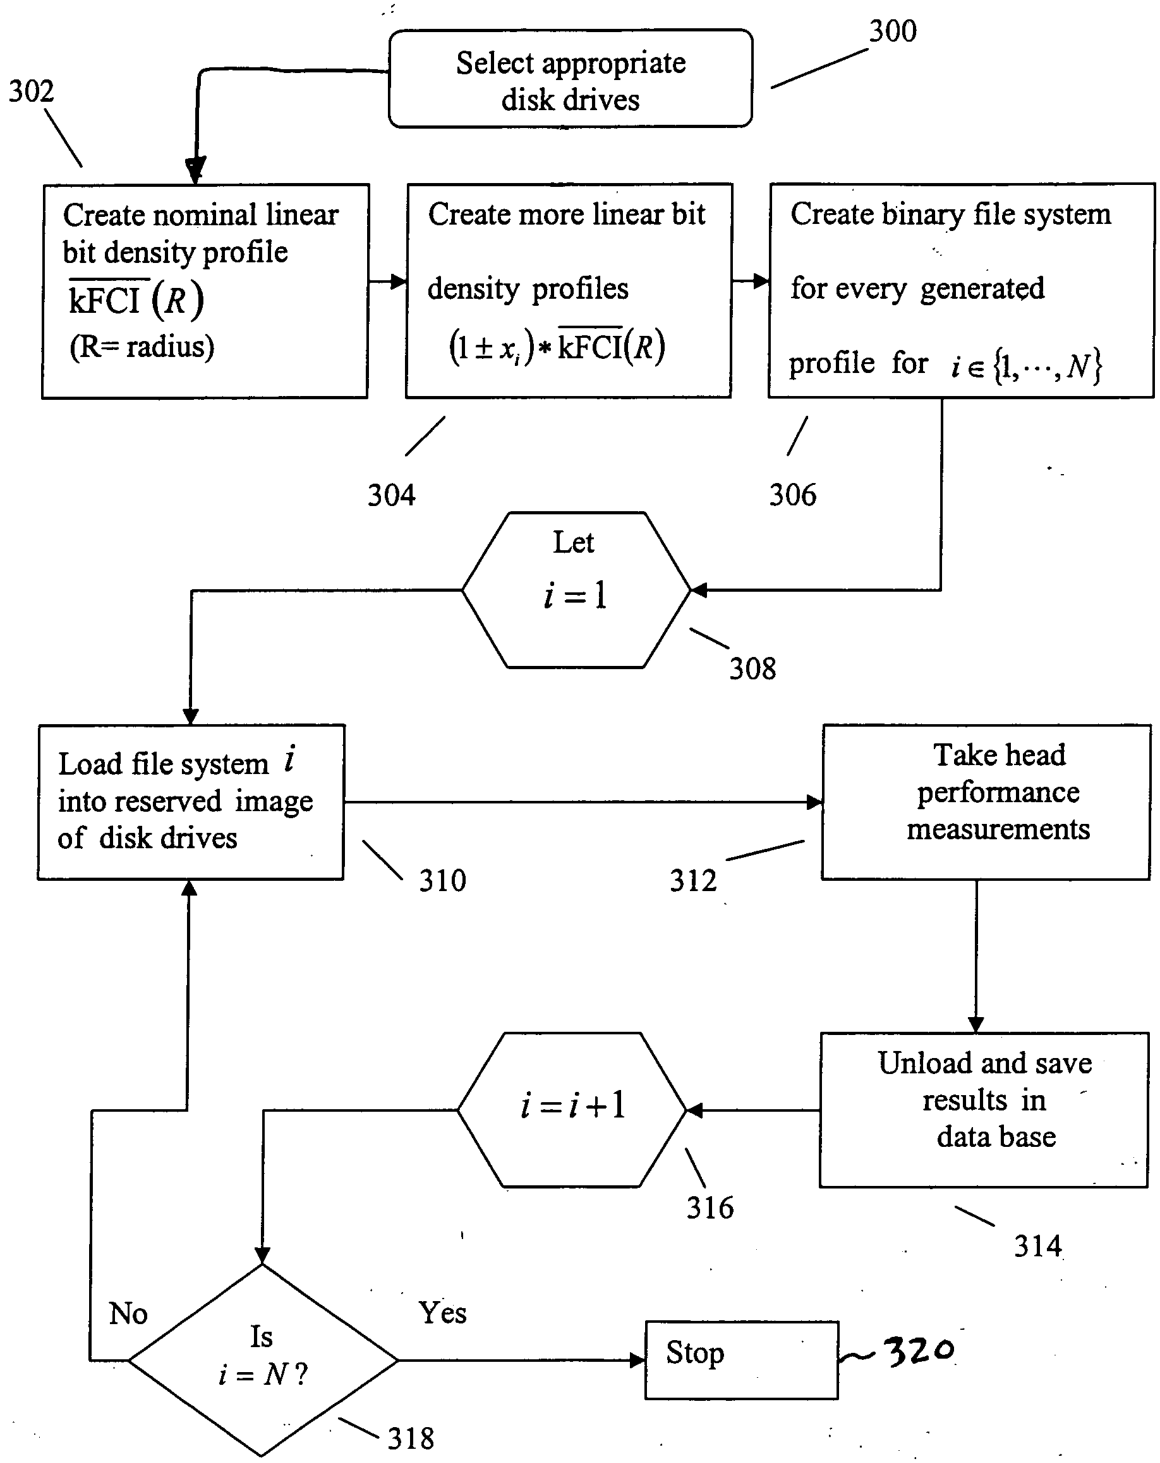
Select appropriate disk drives
300
302
Create nominal linear bit density profile $\overline{\text{kFCI}}(R)$ (R= radius)
Create more linear bit density profiles $(1 \pm x_i) * \overline{\text{kFCI}}(R)$
Create binary file system for every generated profile for $i \in \{1, \cdots, N\}$
304
306
Let $i = 1$
308
Load file system i into reserved image of disk drives
310
312
Take head performance measurements
Unload and save results in data base
$i = i + 1$
316
314
No
Is $i = N$?
Yes
Stop
320
318

FIG. 3

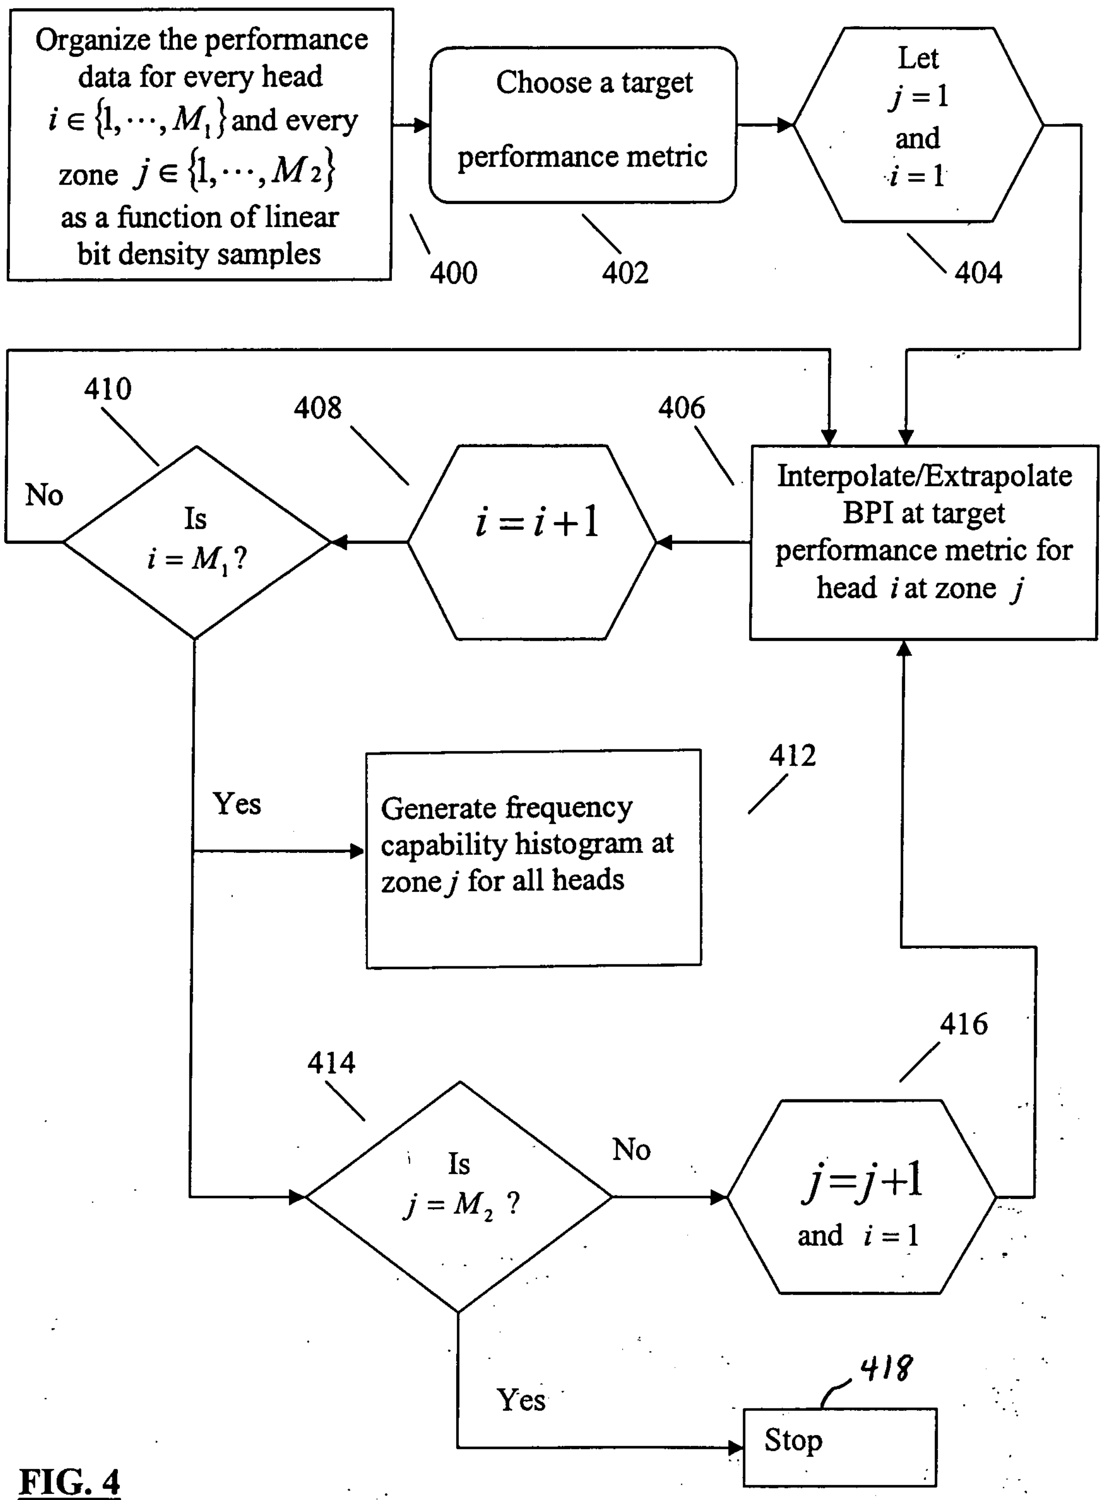
Organize the performance data for every head $i \in \{1, \cdots, M_1\}$ and every zone $j \in \{1, \cdots, M_2\}$ as a function of linear bit density samples
Choose a target performance metric
Let $j=1$ and $i=1$
400
402
404
410
408
406
No
Is $i = M_1$?
$i = i+1$
Interpolate/Extrapolate BPI at target performance metric for head i at zone j
412
Yes
Generate frequency capability histogram at zone j for all heads
416
414
Is $j = M_2$?
No
$j=j+1$ and $i=1$
418
Yes
Stop

FIG. 4

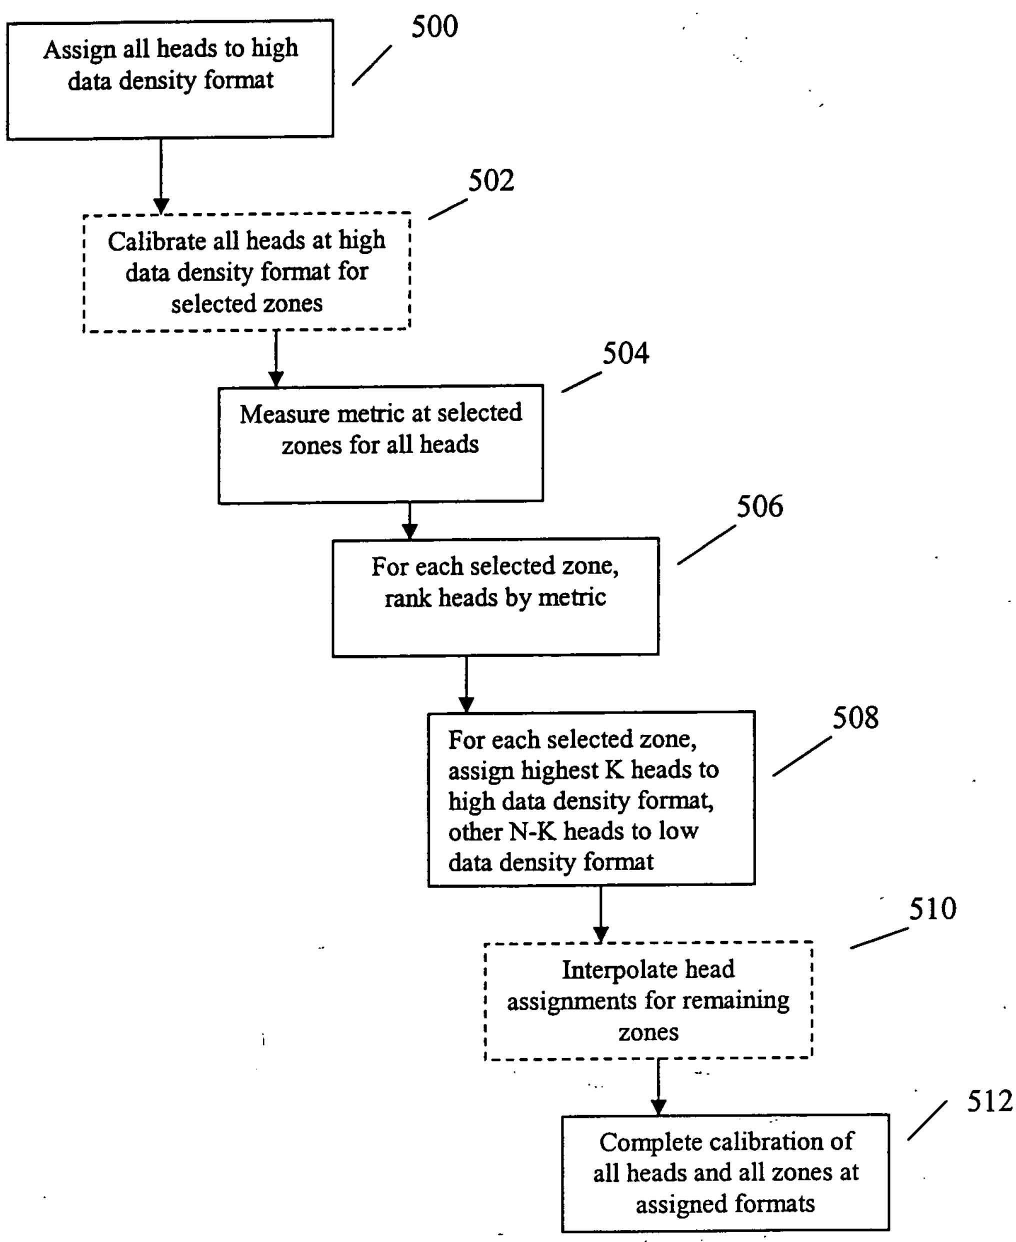
Assign all heads to high data density format
500
Calibrate all heads at high data density format for selected zones
502
Measure metric at selected zones for all heads
504
For each selected zone, rank heads by metric
506
For each selected zone, assign highest K heads to high data density format, other N-K heads to low data density format
508
Interpolate head assignments for remaining zones
510
Complete calibration of all heads and all zones at assigned formats
512

FIG. 5

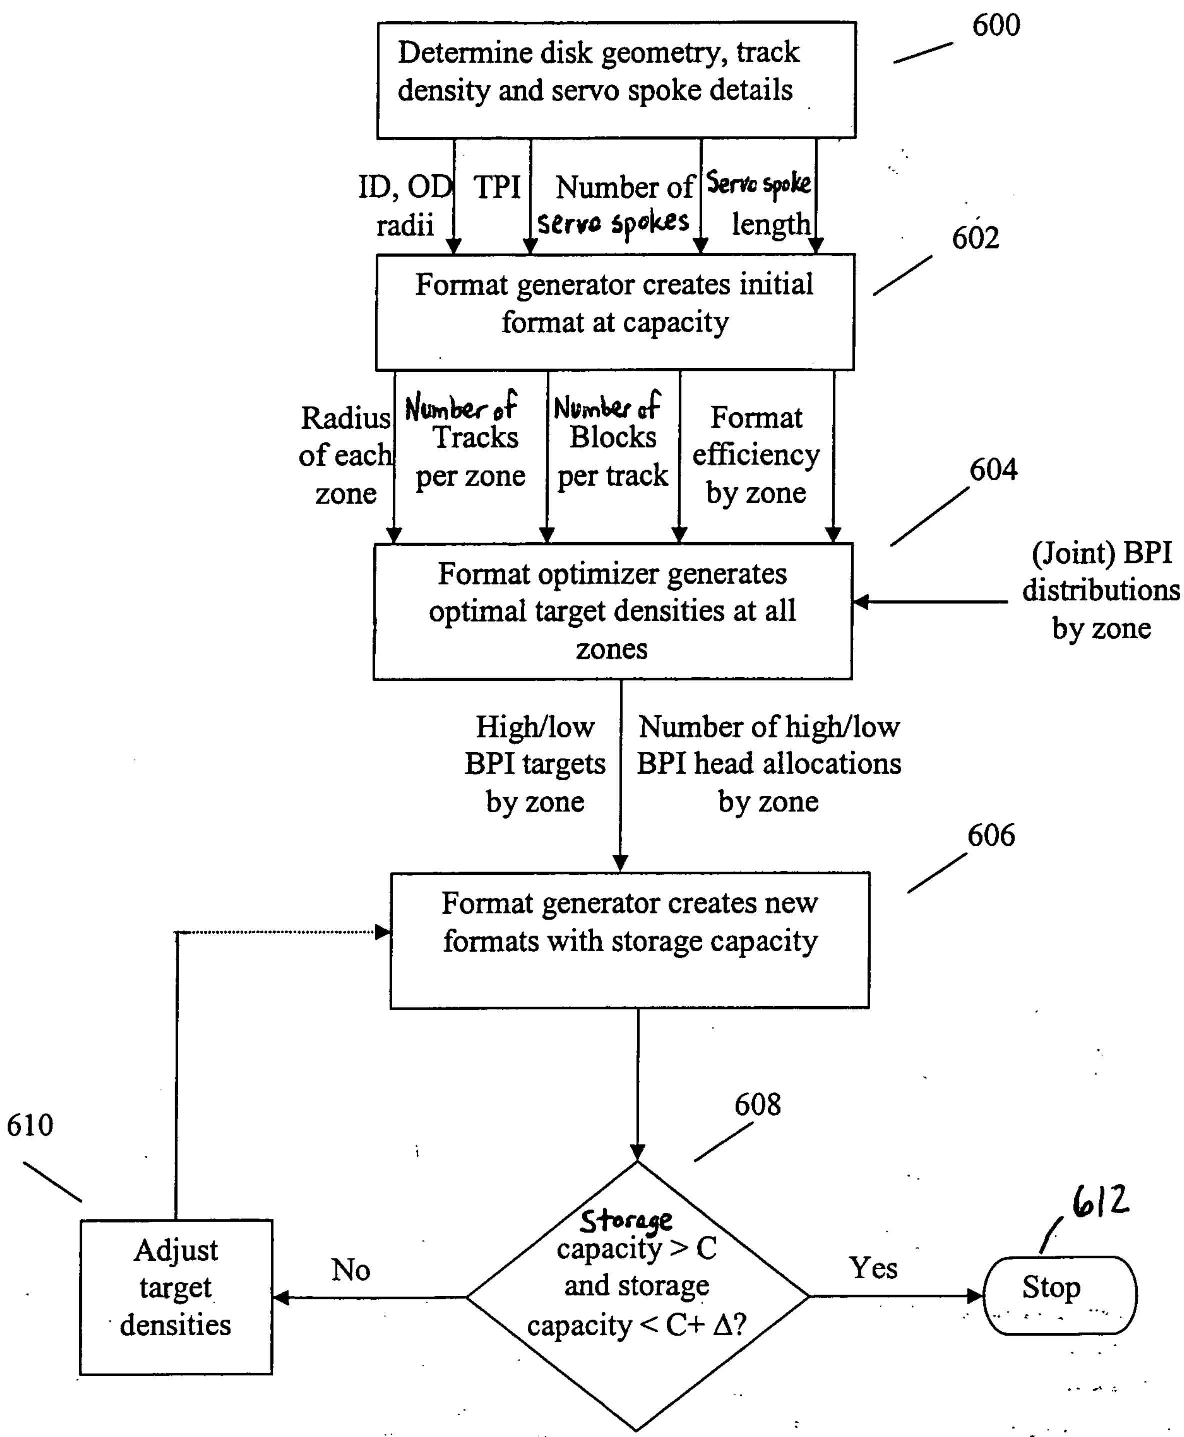
600
Determine disk geometry, track density and servo spoke details
ID, OD radii
TPI
Number of servo spokes
Servo spoke length
602
Format generator creates initial format at capacity
Radius of each zone
Number of Tracks per zone
Number of Blocks per track
Format efficiency by zone
604
Format optimizer generates optimal target densities at all zones
(Joint) BPI distributions by zone
High/low BPI targets by zone
Number of high/low BPI head allocations by zone
606
Format generator creates new formats with storage capacity
608
Storage capacity > C and storage capacity < C+ Δ?
No
Yes
610
Adjust target densities
612
Stop

FIG. 6